IBM PowerVM
FW950.30 and FW1010.10
with VIOS 3.1.3.10 operating on
IBM Power Systems
POWER9 and Power10 hardware

This document contains the following information in support of the EAL 2 certification

- IBM Power 3.1.3 User Guidance
- Virtual I-O Server
- PowerVM 3.1.3 Logical Partitioning
- Setting up the virtualization environment
- Managing the virtualization environment
- Monitoring the virtualization environment
- POWER9 Beginning troubleshooting and problem analysis
- POWER9 Installing and configuring the Hardware Management Console
- POWER9 Managing the Hardware Management Console
- POWER9 Problem analysis system parts and locations for the IBM Power Systems HMC
- Power10 Beginning troubleshooting and problem analysis
- Power10 Installing and configuring the Hardware Management Console
- Power10 Managing the Hardware Management Console
- Power10 Problem analysis system parts and locations for the IBM Power Systems HMC
- Power10 Servicing the IBM Power Systems HMC (7063-CR2)

IBM Power 3.1.3 User Guidance

Revision 1.1

Jun 21st, 2022

Prepared By:
IBM
3605 US Hwy. 52 North
Rochester, MN   55901

1 About This Information

This information will tell you how to plan, install, set up and manage the logical partitioning of your Power server based for Common Criteria evaluation.

The IBM PowerVM Architecture on POWER Systems listed in the IBM PowerVM 3.1.3 with VIOS 3.1.3 for POWER9 and Power10 Security Target. The hardware and firmware allow you to set up more than one virtual machine on your server, so that you can run separate operating systems concurrently. Each virtual machine is called a *partition or logical partition (LPAR)*. The design of the architecture provides the following security features:

- The hardware and firmware provide the operating system on each separate partition with the resources it needs to function
- The hardware and firmware keep the resources for each partition separate, so that they will not interfere with each other.

The IBM Logical Partitioning Architecture on Power has been developed and evaluated in accordance with the Common Criteria EAL2 (Evaluation Assurance Level) assurance requirements listed below:

Objectives

- This certification permits a developer to gain assurance from positive security engineering based on good commercial development practices which, though rigorous, do not require substantial specialist knowledge, skills, and other resources. This is an EAL2 evaluation which is the highest level that is recognized globally by all Common Criteria certification entities.
- This certification is therefore applicable in those circumstances where developers or users require a level of independently assured security in conventional commodity targets of evaluation (TOE) and are prepared to incur additional security specific engineering costs.

Assurance components

- The evaluation provides assurance by an analysis of the security functions, using a functional and complete interface specification, guidance documentation, the high-level and low-level design of the TOE, and a subset of the implementation, to understand the security behavior. Assurance is additionally gained through an informal model of the TOE security policy. The analysis is supported by independent testing of the TOE security functions, evidence of developer testing based on the functional specification and high-level design, selective independent confirmation of the developer test results, strength of function analysis, evidence of a developer search for vulnerabilities, and an independent vulnerability analysis demonstrating resistance to penetration attackers with a low attack potential.
- The evaluation also provides assurance through the use of development environment.

Common Criteria security requires the following documentation:

- *Administrator Guidance*, which describes the tasks that a security administrator must perform to install and manage a Common Criteria-evaluated system.
- *User Guidance*, which describes the user's responsibilities for security. In this case, once a POWER server has been configured to run with multiple partitions, the user of the partition does not need to do anything to support security. All the security features are enforced by the firmware and hardware.

This information is designed to meet the Common Criteria requirement for administrator guidance, when used together with the following documents:

- Information about installing, configuring, managing, and servicing the HMC can all be accessed from the HMC itself. This information is part of the IBM Knowledge Center offering and is available on the internet. For example, information about Managing the HMC can be found in documents *PowerVM 3.1.3 POWER9 Managing the Hardware Management Console.pdf and PowerVM 3.1.3 Power10 Managing the Hardware Management Console.pdf*.
- Information about troubleshooting, service and support of the Hardware Management Console (HMC) can be found in documents:
- PowerVM 3.1.3 POWER9 Servicing the IBM Power Systems HMC (7063-CR2).pdf
- PowerVM 3.1.3 POWER9 Beginning troubleshooting and problem analysis.pdf
- PowerVM 3.1.3 POWER9 Problem analysis system parts and locations for the IBM Power Systems HMC.pdf
- PowerVM 3.1.3 Power10 Servicing the IBM Power Systems HMC (7063-CR2).pdf
- PowerVM 3.1.3 Power10 Beginning troubleshooting and problem analysis.pdf
- PowerVM 3.1.3 Power10 Problem analysis system parts and locations for the IBM Power Systems HMC.pdf

Information about creating a virtual computing environment on Power servers can be found in document *PowerVM 3.1.3 Logical Partitioning.pdf*.
You should read these guides first, and you should consider it your primary source of information for setting up logical partitioning of your POWER server to meet the Common Criteria security requirements that are listed in the IBM PowerVM 3.1.3 with VIOS 3.1.3 for POWER9 and Power10 Security Target.

2 Who should read this information

This information is intended for system administrators or security administrators that want to customize a Power server with Logical Partitioning within the valid Common Criteria configuration. This information details the unique requirements of Common Criteria security, and it is intended as a supplement to other manuals describing how you install and set up your system.

3 Overview of security features

There are three categories of security features provided by the IBM Logical Partition Architecture implementation for Power:

1. The Logical Partitioning Architecture implementation ensures that resources can be assigned to partitions by an authorized user and that those resources will not be accessible to other partitions.

2. The Logical Partitioning Architecture implementation ensures that communication between partitions can occur only using channels established by an authorized user.

3. The Logical Partitioning Architecture implementation ensures that each partition cannot access resources or communicate with other partitions except when explicitly allowed by an authorized user.

In addition, the following assumptions are made about the operating environment when the system is in operation:

1. A suitable management console must be configured for use by a capable and trustworthy user assigned to follow the applicable guidance in order to install and operate the system within the evaluated configuration.

2. The system must be installed and configured in accordance with its guidance documents, including connecting appropriate device.

3. The system must be within a physical environment suitable to protect itself and its external connections from inappropriate access and modification.

4 Security Target of Evaluation (TOE)

The IBM PowerVM 3.1.3 with VIOS 3.1.3 for POWER9 and Power10 Security Target of Evaluation (TOE) is the combination of hardware and software that provides security protection within a computer system. The TOE:

- POWER9 E980 System with firmware level 01VH950_092_045
- Power10 E1080 System with firmware level 01MH1010_094_094

POWER servers can be configured to house multiple independent systems within the same server. Each independent system within the server is called a partition. The partitions do not have to run the same type of operating system. During the configuration process, an administrator determines what resources within the server will be assigned to each independent partitions. There are many partition features. The following Logical Partitioning Architecture features are allowed in the evaluated configuration:

- Micro-partitioning: This feature allows a processor to be shared between two partitions. One partition may get 10% while another gets 90%.

- Virtual Ethernet: This feature provides optional communication between partitions on the server. A virtual network switch is also supported to control the traffic.

- The Logical Partitioning Architecture was evaluated independent of the OS in the partition. Any operating system may be installed in the partition.

Products that are included in the TOE have been evaluated and tested for Common Criteria security compliance. Products that are not included in the TOE have not been evaluated for Common Criteria security compliance. Because the TOE is a general building block, many installations require changes or additions to the evaluated configuration. Your security administrator should assess the security risk of any changes or additions to the TOE configuration.

4.1 Physical System security

To be within the evaluated configuration, the system must be in a secured room with limited and monitored access. The systems HMC appliance must also be in the same secured area.

4.2 System Management Console installation

A common criteria compliant system must be configured by a Management Console (MC). It doesn't matter if the MC is directly connected to the system or connected to the system through a network. The MC is used to partition the system.

The guidance for connecting an HMC to a system can be found in documents *PowerVM 3.1.3 POWER9 Installing and configuring the Hardware Management Console.pdf and PowerVM 3.1.3 Power10 Installing and configuring the Hardware Management Console.pdf*.

4.3 System firmware installation

For POWER9 the evaluated firmware level is 01VH950_092_045 which has a sha256sum of aea91c2bd13a00801003e27438cd268b3c31982d57d444187c278deaea0c2714. For Power10 the evaluated firmware level is 01MH1010_094_094 which has a sha256sum of f7828e0b46f9fa4a3326707c0af25401304deb83ef7ff05c83cbf24cb37490f8. Verify this level of firmware is on your machine. To check the level of firmware on your system, follow these instructions:

1. From the HMC, In the navigation area, click the Resources icon, and then select All Servers.
2. Select the server for which you want to view system information.
3. In the menu pod, expand Actions and then expand Updates.
4. Select View system information
5. In the Specify LIC Repository window, select None – Display current values and click ok.

If the firmware level does not match the evaluated firmware level, you must install the evaluated firmware on your machine. The firmware is available for download from IBM fix central https://www.ibm.com/support/fixcentral/ . Enter your product 9080-M9S (POWER9) or 9080-HEX (Power10), specify the base firmware release, VH950 (POWER9) or NH1010 (Power10) and specify the specific certified firmware level.

If you need to install the evaluated firmware on your machine, follow these instructions:

1. From the HMC, In the navigation area, click the Resources icon, and then select All Servers.
2. Select the server for which you want to update system information and click Actions > Updates.
3. Select Change Licensed Internal Code > for the Current Release or to a new Release based on the currently installed release.
4. Select an action from the list and click Ok.
5. When you complete this task, click Close.

4.4 VIOS installation

The VIOS level that will be used for the evaluation is VIOS 3.1.3.10. The flash image name is: **Virtual_IO_Server_Base_Install_3.1.3.10_Flash_092021_LCD8250308.iso.** The sha256sum for the above flash image is: **92b4f3af830254861e8d5907f61550ca0a3c222a3e3ee604b9d76a49a26fbeaa**

This VIOS level is ordered and entitled with the purchase of a Power Server. The detailed steps to download the above image are given below:

Here are the download instructions using for the PRPQ.

1) Go to the IBM site and navigate to the ESS site.

a) Go to https://www.ibm.com *(See Figure 1 in Appendix)*

b) Click on "**Learn & Support**" pull down menu

c) Under "Support" click on "**View more on Support**" *(See Figure 2 in Appendix)*

d) Scroll down to the "**Downloads, fixes & updates**" section *(See Figure 3 in Appendix*

e) Click on "**IBM Power & Storage**", *(See Figure 4 in Appendix)*

f) Scroll down and click on "**Sign in**", to sign in with your IBM Web ID, *(See Figure 5 in Appendix)*

2) If you have never been on the site before you must "attach" yourself to your IBM Customer Number.

a) Click on "**My profile**". *(See Figure 6 in Appendix)*

b) Click on "**Register customer number**".*(See Figure 7 in Appendix)*

c) You may enter the country code/customer number combination or the HW/SW serial number of an IBM product purchased using that customer number. IF you are the first to register for the customer, you will become the "primary" contact for that customer number and have to approve future requests to attach Web IDs to that customer number. If you are not the first, your request to attach the customer number will be sent to the current primary contact for that customer number. You can go no further until your IBM Web ID is associated with one or more customer numbers.

3) Start the software download. *(See Figure 5 in Appendix)*

a) Click on "**My entitled software**". *(See Figure 8 in Appendix)*
b) Click on "**Software download**s". *(See Figure 9 in Appendix)*
c) At this point, there are several different ways to look at your entitled software. Since this note is specifically for the Common Criteria certification of VIOS 3.1.3.10, I will use the "By product" tab. Click on the "**By product**" tab. *(See Figure 10 in Appendix)*

d) Find the product, 5799-P31 (PowerVM V3 Technology Level), then press "**Add product**". *(See Figure 11 in Appendix)*

e) Click on the "**Continue**" button. You will get your selected product and the delivery features under that product. *(See Figure 12 in Appendix)*
f) Click on the "**packages**" button next to "**6018: VIO 3.1.3.10 v03.01.03,ENU,DVD** ". *(See Figure 13 in Appendix)*
g) Click the check box next to the "**ISO, Virtual I/O Server v3.1.3.10 Flash (9/2021)**" *(See Figure 14 in Appendix)*

h) Having selected the file you want to download, click the "**Continue**" button. *(See Figure 15 in Appendix)*

i) Click on the "**I agree**" button to accept the licenses. *(See Figure 16 in Appendix)*
j) Choose your download method either HTTPS (or Download Director) and click the "**Continue**" button. *(See Figure 17 in Appendix)*

k) Click on the "**Virtual_IO_Server_Base_Install_3.1.3.10_Flash_092021_LCD8250308.iso** " link to begin the download of the file.

4.5 Partition the Server:

Additional information about how to setup, manage and monitor the virtualization environment can be found in the following documents:

- *PowerVM 3.1.3 Setting up the virtualization environment.pdf*
- *PowerVM 3.1.3 Managing the virtualization environment.pdf*
- *PowerVM 3.1.3 Monitoring the virtualization environment.pdf*

The information on configuring and installing the Virtual I/O Server can be found in document *PowerVM 3.1.3 Virtual I-O Server.pdf.*

There are no special instructions for partitioning the server. However, some features are not allowed in the evaluated configuration. The following features are excluded from the evaluated configuration:

- I/O Pools

  - You must not create a storage pool, and you must remove any storage pools that exist. To check if you have any storage pools, do the following:
    - From the HMC, select properties for a partition.
    - Follow the tabs hardware -> I/O.
    - Press the I/O Pools button
    - If any storage pools exist, you must remove them. Also, you must not add storage pools from this screen at any time.

    Note: If any storage pools were removed, the partition must be rebooted.
- Workload Management Groups
  - Partitions must be configured with a workload management group of none (i.e. a numbered group is not supported. To check if you have a workload management group configured, do the following:
    - From the HMC, select Configuration-> Manage Profiles for a partition.
    - Select the partition's and choice the action profile action.
    - Select the Actions -> Edit.
    - Inspect the profile and ensure workload management group is none
- Power Controlling
  - You must not allow any partition to have a power controlling partition, and you must remove any power controlling partitions if they exist. To check that a partition does not have a power controlling partition, do the following:
    - From the HMC, select Configuration-> Manage Profiles for a partition.
    - Select the partition's profile check box.
    - Select the Actions -> Edit.
    - Select the Power Controlling tab.
    - If there are any partitions in the power controlling partitions list, select them and press the remove button.

Note: If any controlling partitions were removed, the partition must be rebooted.

- Shared Storage Pools
    - You must not create a shared storage pool. This feature is supported, but not activated on a freshly installed Virtual I-O Server.
- Cache Management
    - You must not enable cache management. This feature is disabled on a freshly installed Virtual I-O Server.

While the Virtual I/O Server supports various virtualization features, the scope of this evaluation only includes Virtual SCSI, Virtual Ethernet and Shared Ethernet Adapter (SEA) technologies. NPIV is not covered in this evaluation as much of the storage administration and isolation enforcement in the NPIV framework is provided in the SAN infrastructure, external to PowerVM and the server.

NOTE: The selection and installation of the individual operating systems for the partitions is outside the scope of this evaluation. The evaluated hardware and firmware are indifferent to the OS of the partition.

4.6 Additional Information about the PowerVM Hypervisor, VIOS and HMC

The following is some additional information/considerations about the environment of PowerVM and VIOS:

- The PowerVM Hypervisor is unlike some other products that have a normal mode of operation and a diagnostic mode of operation. When booting the PowerVM Hypervisor, there is only a normal mode of operation and does not support a separate diagnostic boot.
- The POWER hardware platform only supports a single version level for a given hardware platform. POWER9 systems can only run firmware written specifically for POWER9 hardware. Similarly, Power10 hardware can only run firmware specifically written for Power10 hardware. The VIOS is written such that the same VIOS level can be run on different hardware platforms. For example, PowerVM 3.1.3 version of the VIOS is supported both on POWER9 and Power10 hardware platforms.
- The HMC is a separate entity running that is not running on the POWER server under test. A single HMC has the ability to manage different versions of POWER servers simultaneously. A single HMC with version 10 of the HMC image could be managing Power10, POWER9 and POWER8 servers.

APPENDIX

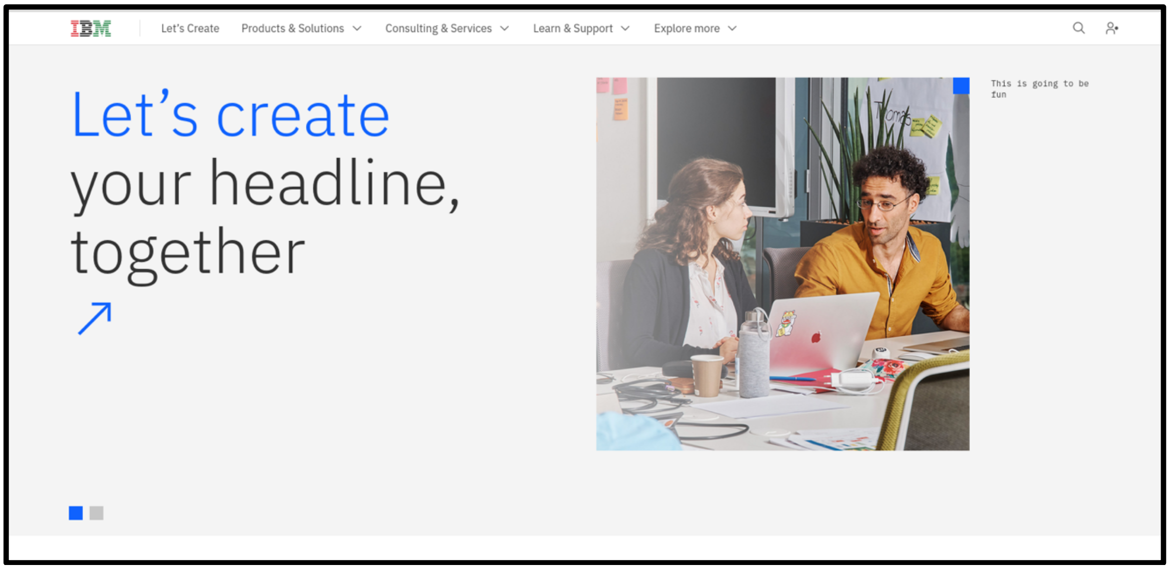

Let's create your headline, together

This is going to be fun

Figure 1: IBM Site

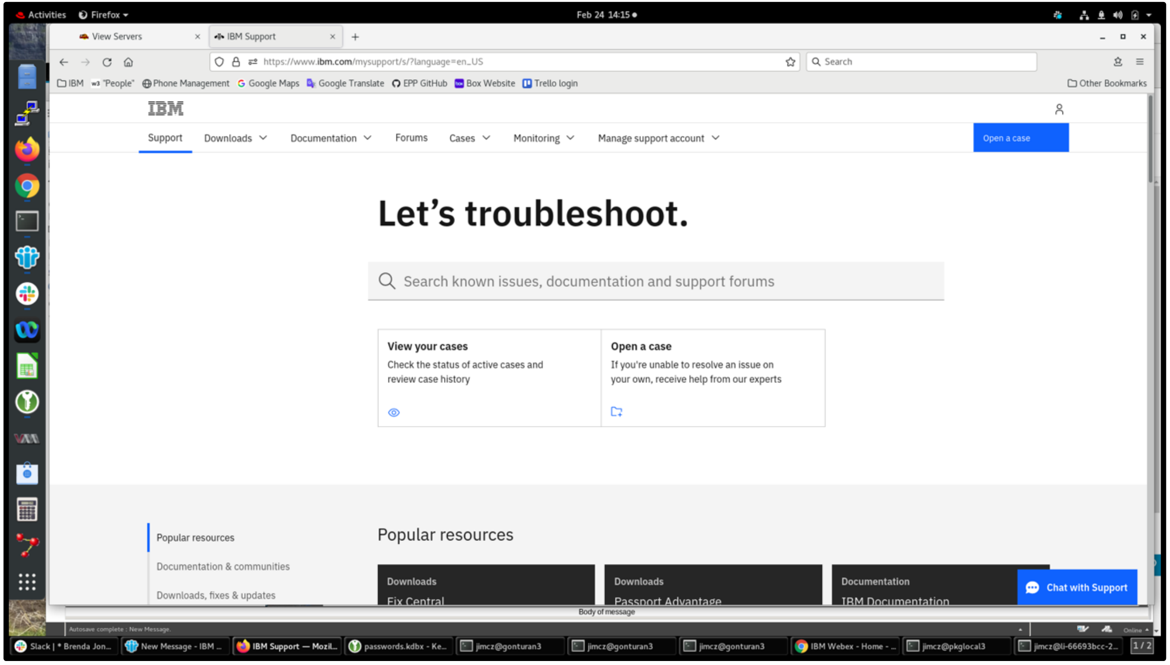

Let's troubleshoot.

Search known issues, documentation and support forums

View your cases
Check the status of active cases and review case history

Open a case
If you're unable to resolve an issue on your own, receive help from our experts

Popular resources

Figure 2: Support Page

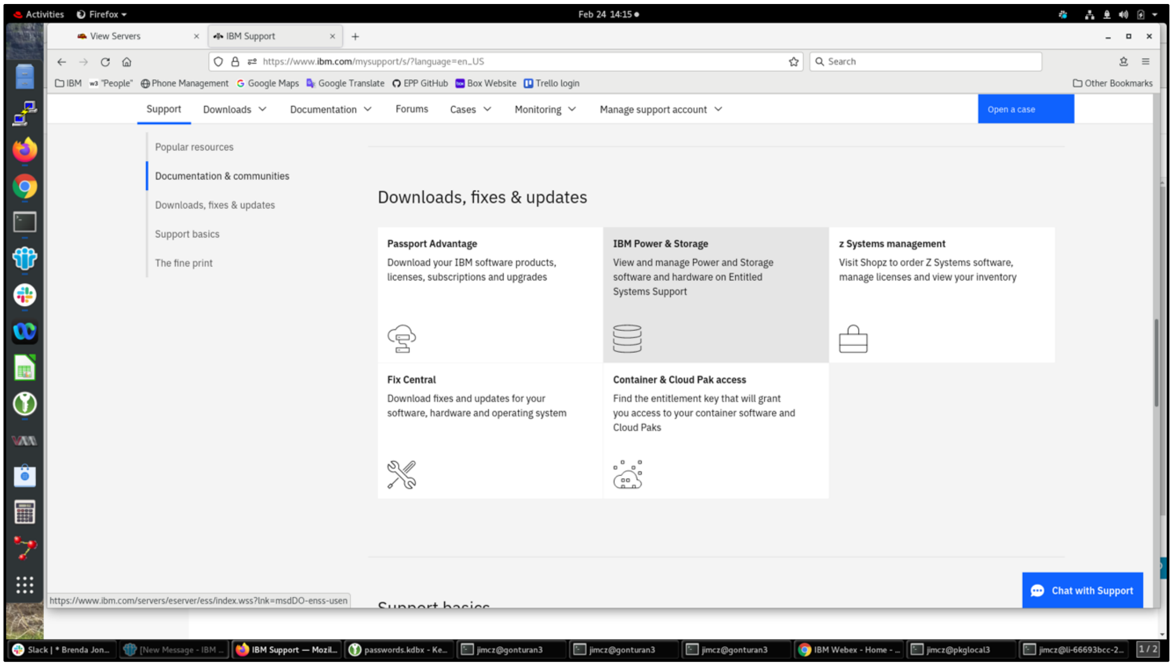

Downloads, fixes & updates

Passport Advantage
Download your IBM software products, licenses, subscriptions and upgrades

IBM Power & Storage
View and manage Power and Storage software and hardware on Entitled Systems Support

z Systems management
Visit Shopz to order Z Systems software, manage licenses and view your inventory

Fix Central
Download fixes and updates for your software, hardware and operating system

Container & Cloud Pak access
Find the entitlement key that will grant you access to your container software and Cloud Paks

Figure 3: Downloads, Fixes and Updates

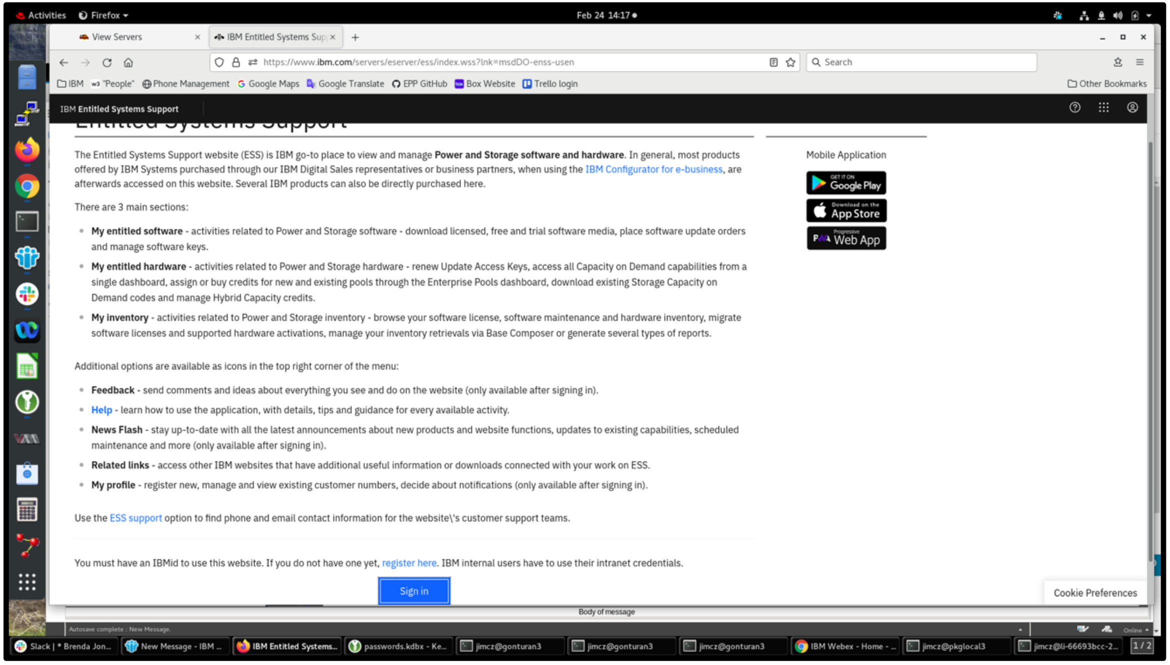

The Entitled Systems Support website (ESS) is IBM go-to place to view and manage **Power and Storage software and hardware**. In general, most products offered by IBM Systems purchased through our IBM Digital Sales representatives or business partners, when using the IBM Configurator for e-business, are afterwards accessed on this website. Several IBM products can also be directly purchased here.

There are 3 main sections:

- **My entitled software** - activities related to Power and Storage software - download licensed, free and trial software media, place software update orders and manage software keys.
- **My entitled hardware** - activities related to Power and Storage hardware - renew Update Access Keys, access all Capacity on Demand capabilities from a single dashboard, assign or buy credits for new and existing pools through the Enterprise Pools dashboard, download existing Storage Capacity on Demand codes and manage Hybrid Capacity credits.
- **My inventory** - activities related to Power and Storage inventory - browse your software license, software maintenance and hardware inventory, migrate software licenses and supported hardware activations, manage your inventory retrievals via Base Composer or generate several types of reports.

Additional options are available as icons in the top right corner of the menu:

- **Feedback** - send comments and ideas about everything you see and do on the website (only available after signing in).
- **Help** - learn how to use the application, with details, tips and guidance for every available activity.
- **News Flash** - stay up-to-date with all the latest announcements about new products and website functions, updates to existing capabilities, scheduled maintenance and more (only available after signing in).
- **Related links** - access other IBM websites that have additional useful information or downloads connected with your work on ESS.
- **My profile** - register new, manage and view existing customer numbers, decide about notifications (only available after signing in).

Use the ESS support option to find phone and email contact information for the website\'s customer support teams.

You must have an IBMid to use this website. If you do not have one yet, register here. IBM internal users have to use their intranet credentials.

Figure 4: ESS Entry

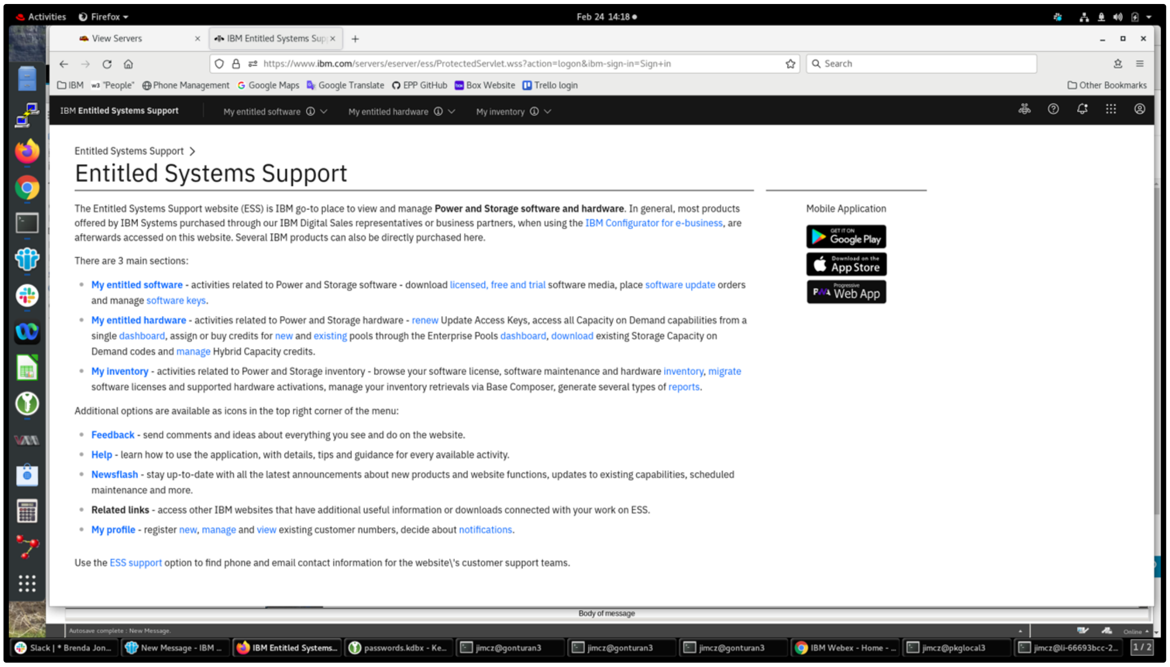

Entitled Systems Support >

Entitled Systems Support

The Entitled Systems Support website (ESS) is IBM go-to place to view and manage **Power and Storage software and hardware**. In general, most products offered by IBM Systems purchased through our IBM Digital Sales representatives or business partners, when using the IBM Configurator for e-business, are afterwards accessed on this website. Several IBM products can also be directly purchased here.

There are 3 main sections:

- **My entitled software** - activities related to Power and Storage software - download licensed, free and trial software media, place software update orders and manage software keys.
- **My entitled hardware** - activities related to Power and Storage hardware - renew Update Access Keys, access all Capacity on Demand capabilities from a single dashboard, assign or buy credits for new and existing pools through the Enterprise Pools dashboard, download existing Storage Capacity on Demand codes and manage Hybrid Capacity credits.
- **My inventory** - activities related to Power and Storage inventory - browse your software license, software maintenance and hardware inventory, migrate software licenses and supported hardware activations, manage your inventory retrievals via Base Composer, generate several types of reports.

Additional options are available as icons in the top right corner of the menu:

- **Feedback** - send comments and ideas about everything you see and do on the website.
- **Help** - learn how to use the application, with details, tips and guidance for every available activity.
- **Newsflash** - stay up-to-date with all the latest announcements about new products and website functions, updates to existing capabilities, scheduled maintenance and more.
- **Related links** - access other IBM websites that have additional useful information or downloads connected with your work on ESS.
- **My profile** - register new, manage and view existing customer numbers, decide about notifications.

Use the ESS support option to find phone and email contact information for the website\'s customer support teams.

Figure 5: ESS Main

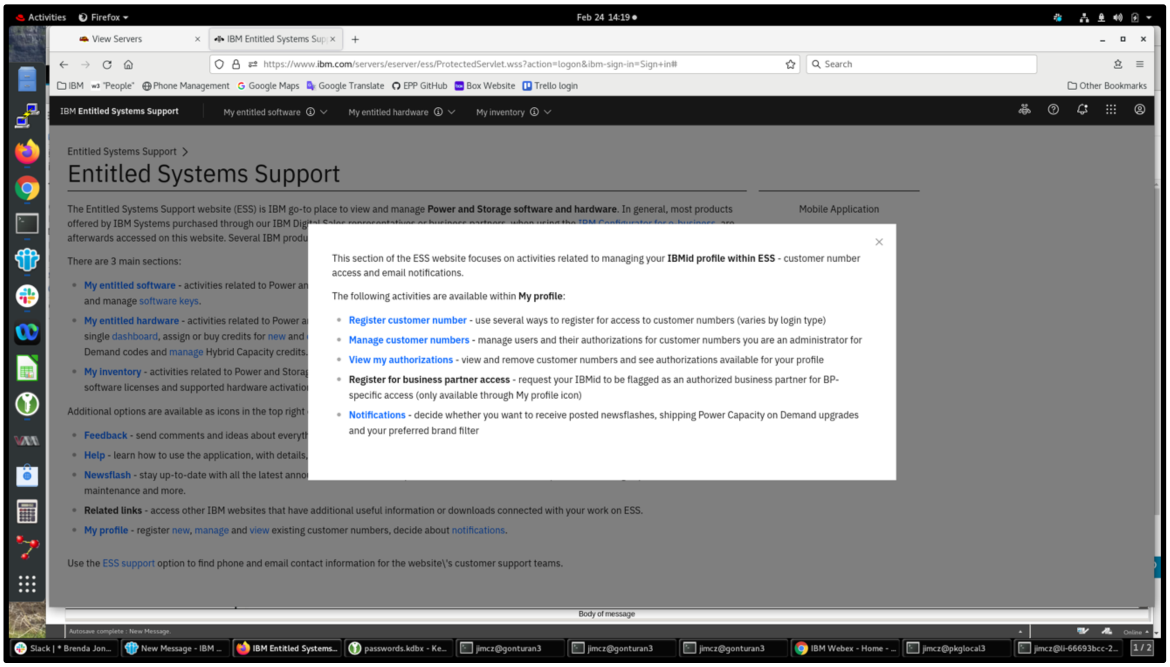

This section of the ESS website focuses on activities related to managing your **IBMid profile within ESS** - customer number access and email notifications.

The following activities are available within **My profile**:

- **Register customer number** - use several ways to register for access to customer numbers (varies by login type)
- **Manage customer numbers** - manage users and their authorizations for customer numbers you are an administrator for
- **View my authorizations** - view and remove customer numbers and see authorizations available for your profile
- **Register for business partner access** - request your IBMid to be flagged as an authorized business partner for BP-specific access (only available through My profile icon)
- **Notifications** - decide whether you want to receive posted newsflashes, shipping Power Capacity on Demand upgrades and your preferred brand filter

Figure 6: My Profile

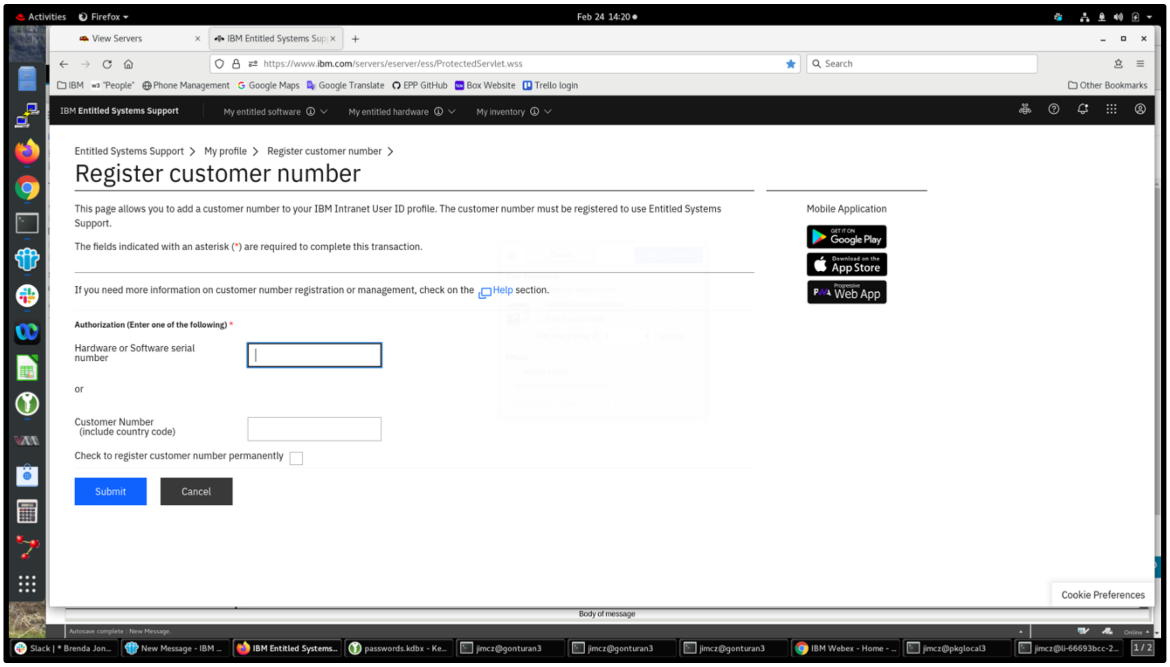

Entitled Systems Support > My profile > Register customer number >

Register customer number

This page allows you to add a customer number to your IBM Intranet User ID profile. The customer number must be registered to use Entitled Systems Support.

The fields indicated with an asterisk (*) are required to complete this transaction.

If you need more information on customer number registration or management, check on the Help section.

Authorization (Enter one of the following) *

Hardware or Software serial number

or

Customer Number
(include country code)

Check to register customer number permanently □

Submit   Cancel

Figure 7: Register Customer Number

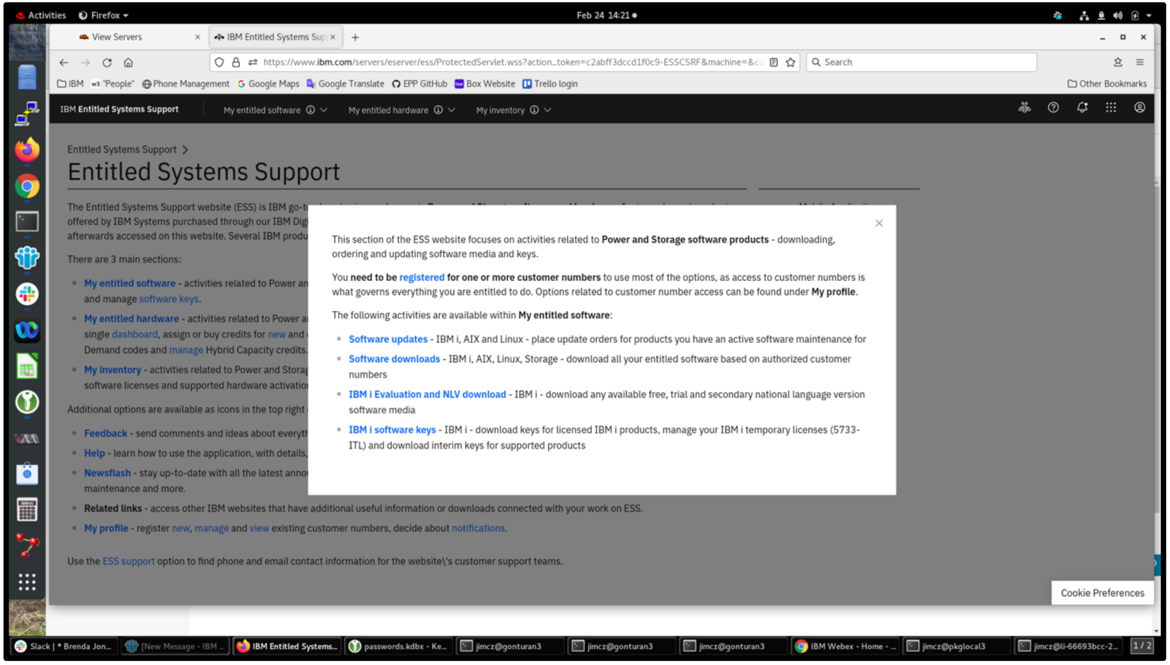
IBM Entitled Systems Support
My entitled software
My entitled hardware
My inventory
https://www.ibm.com/servers/eserver/ess/ProtectedServlet.wss?action_token=c2abff3dccd1f0c9-ESSCSRF&machine=&cu
Cookie Preferences

Entitled Systems Support >

Entitled Systems Support

The Entitled Systems Support website (ESS) is IBM go-t
offered by IBM Systems purchased through our IBM Digi
afterwards accessed on this website. Several IBM produ

There are 3 main sections:

- **My entitled software** - activities related to Power an
  and manage software keys.
- **My entitled hardware** - activities related to Power a
  single dashboard, assign or buy credits for new and
  Demand codes and manage Hybrid Capacity credits.
- **My inventory** - activities related to Power and Storag
  software licenses and supported hardware activatio

Additional options are available as icons in the top right

- **Feedback** - send comments and ideas about everyt
- **Help** - learn how to use the application, with details,
- **Newsflash** - stay up-to-date with all the latest anno
  maintenance and more.
- **Related links** - access other IBM websites that have additional useful information or downloads connected with your work on ESS.
- **My profile** - register new, manage and view existing customer numbers, decide about notifications.

Use the ESS support option to find phone and email contact information for the website\'s customer support teams.

This section of the ESS website focuses on activities related to **Power and Storage software products** - downloading, ordering and updating software media and keys.

You **need to be** registered **for one or more customer numbers** to use most of the options, as access to customer numbers is what governs everything you are entitled to do. Options related to customer number access can be found under **My profile**.

The following activities are available within **My entitled software**:

- **Software updates** - IBM i, AIX and Linux - place update orders for products you have an active software maintenance for
- **Software downloads** - IBM i, AIX, Linux, Storage - download all your entitled software based on authorized customer numbers
- **IBM i Evaluation and NLV download** - IBM i - download any available free, trial and secondary national language version software media
- **IBM i software keys** - IBM i - download keys for licensed IBM i products, manage your IBM i temporary licenses (5733-ITL) and download interim keys for supported products

Figure 8: My Entitled Software

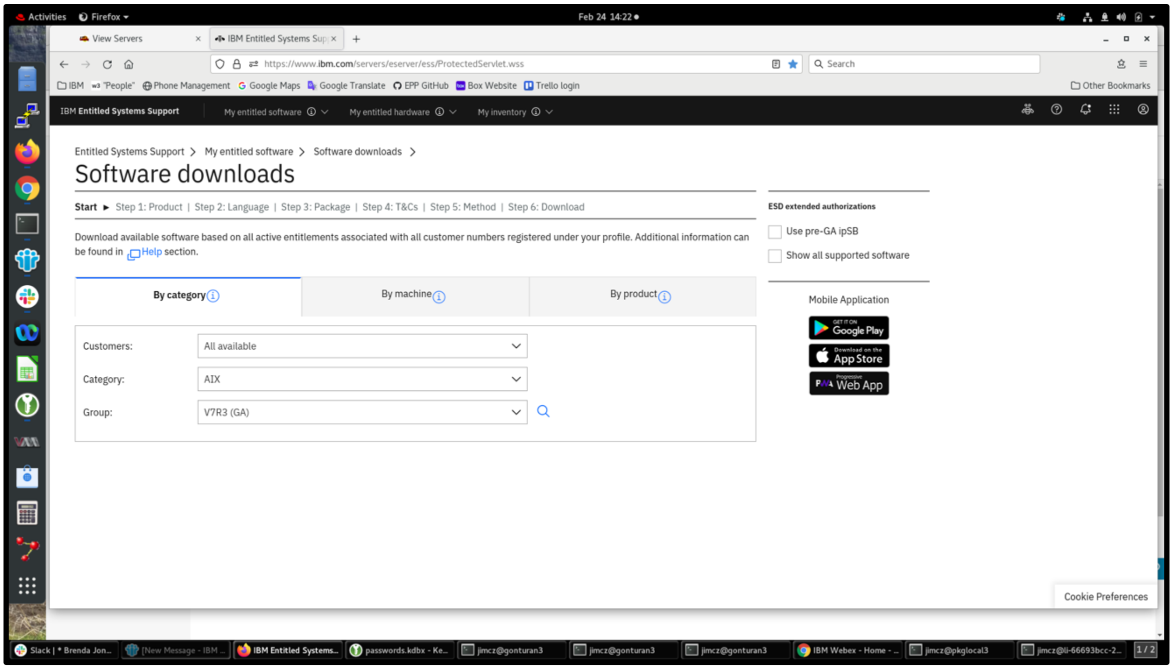

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start ▶ Step 1: Product | Step 2: Language | Step 3: Package | Step 4: T&Cs | Step 5: Method | Step 6: Download

Download available software based on all active entitlements associated with all customer numbers registered under your profile. Additional information can be found in Help section.

By category | By machine | By product

Customers: All available

Category: AIX

Group: V7R3 (GA)

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Mobile Application

Google Play | App Store | Web App

Figure 9: Software Downloads

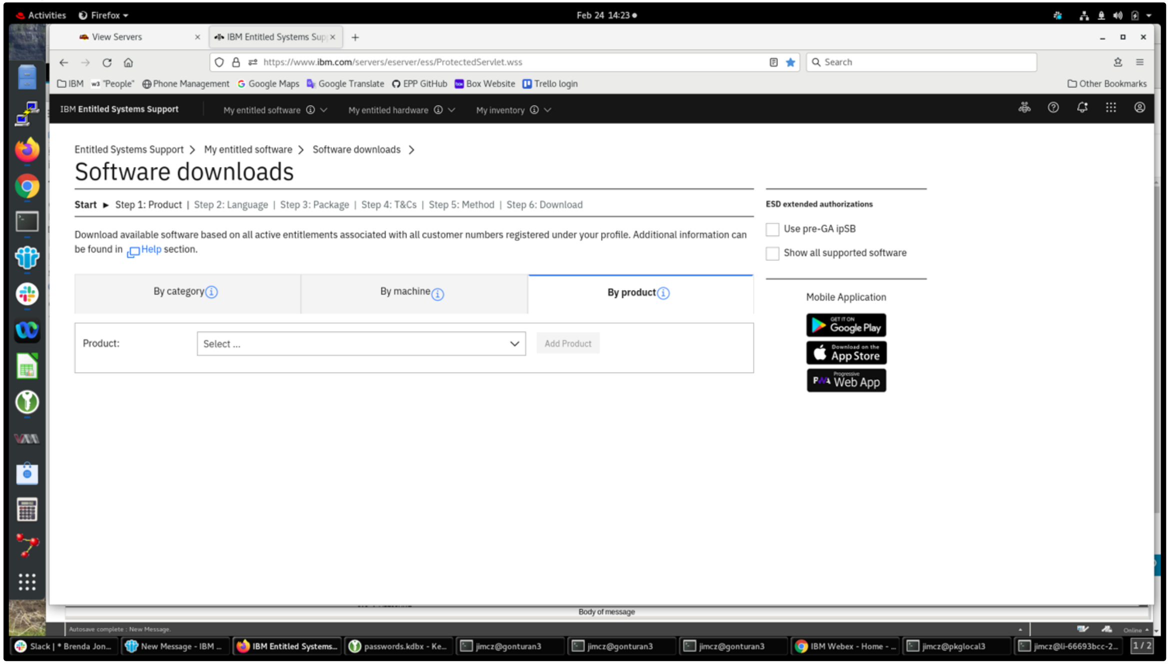

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start ▶ Step 1: Product | Step 2: Language | Step 3: Package | Step 4: T&Cs | Step 5: Method | Step 6: Download

Download available software based on all active entitlements associated with all customer numbers registered under your profile. Additional information can be found in Help section.

By category | By machine | By product

Product: Select ... | Add Product

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Mobile Application

Google Play | App Store | Web App

Figure 10: By Product Tab

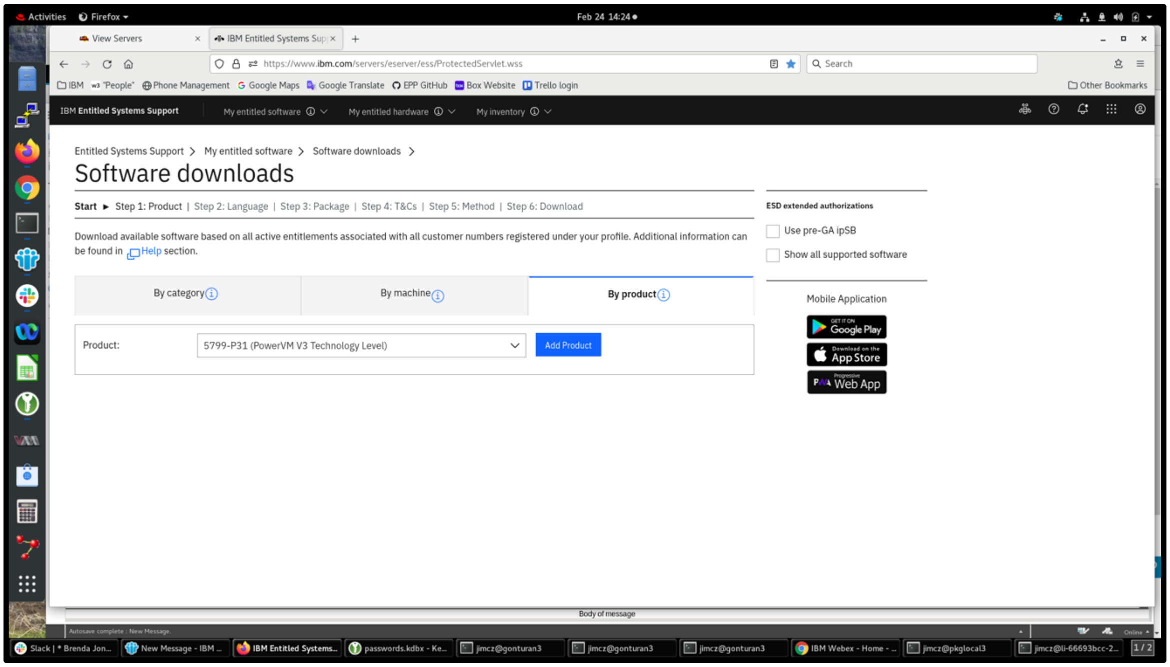

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start ▸ Step 1: Product | Step 2: Language | Step 3: Package | Step 4: T&Cs | Step 5: Method | Step 6: Download

Download available software based on all active entitlements associated with all customer numbers registered under your profile. Additional information can be found in Help section.

By category | By machine | **By product**

Product: 5799-P31 (PowerVM V3 Technology Level) [Add Product]

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Mobile Application

Figure 11: Add 5799-P31

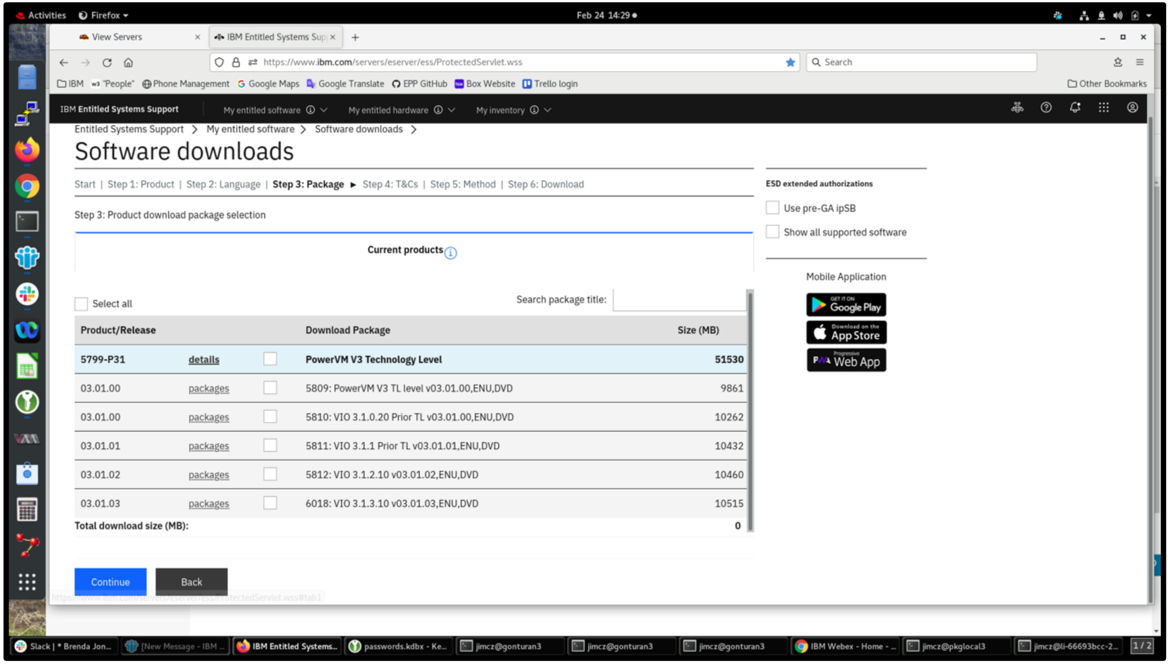

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start | Step 1: Product | Step 2: Language | **Step 3: Package** ▸ Step 4: T&Cs | Step 5: Method | Step 6: Download

Step 3: Product download package selection

Current products

□ Select all

Search package title:

Product/Release			**Download Package**	**Size (MB)**
5799-P31	details	□	**PowerVM V3 Technology Level**	**51530**
03.01.00	packages	□	5809: PowerVM V3 TL level v03.01.00,ENU,DVD	9861
03.01.00	packages	□	5810: VIO 3.1.0.20 Prior TL v03.01.00,ENU,DVD	10262
03.01.01	packages	□	5811: VIO 3.1.1 Prior TL v03.01.01,ENU,DVD	10432
03.01.02	packages	□	5812: VIO 3.1.2.10 v03.01.02,ENU,DVD	10460
03.01.03	packages	□	6018: VIO 3.1.3.10 v03.01.03,ENU,DVD	10515

Total download size (MB): 0

[Continue] [Back]

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Mobile Application

Figure 12: Product Features

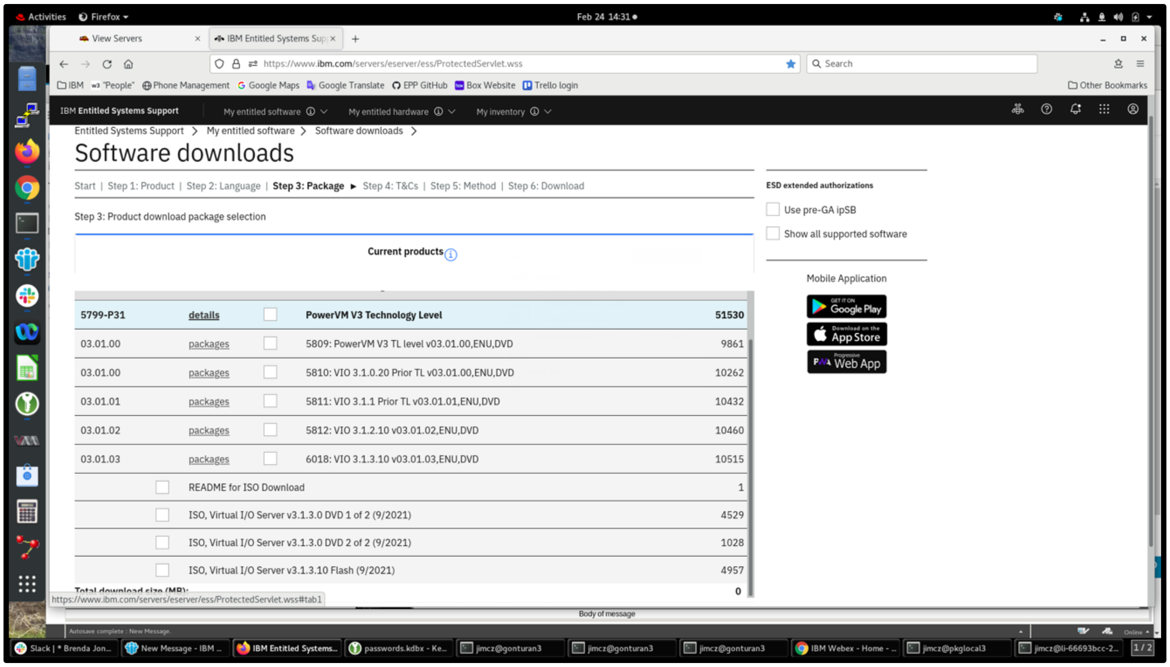

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start | Step 1: Product | Step 2: Language | **Step 3: Package** ▸ Step 4: T&Cs | Step 5: Method | Step 6: Download

Step 3: Product download package selection

Current products

5799-P31	details	□	**PowerVM V3 Technology Level**	**51530**
03.01.00	packages	□	5809: PowerVM V3 TL level v03.01.00,ENU,DVD	9861
03.01.00	packages	□	5810: VIO 3.1.0.20 Prior TL v03.01.00,ENU,DVD	10262
03.01.01	packages	□	5811: VIO 3.1.1 Prior TL v03.01.01,ENU,DVD	10432
03.01.02	packages	□	5812: VIO 3.1.2.10 v03.01.02,ENU,DVD	10460
03.01.03	packages	□	6018: VIO 3.1.3.10 v03.01.03,ENU,DVD	10515
	□	README for ISO Download		1
	□	ISO, Virtual I/O Server v3.1.3.0 DVD 1 of 2 (9/2021)		4529
	□	ISO, Virtual I/O Server v3.1.3.0 DVD 2 of 2 (9/2021)		1028
	□	ISO, Virtual I/O Server v3.1.3.10 Flash (9/2021)		4957

Total download size (MB): **0**

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Mobile Application

Figure 13: 5799-P31 6018

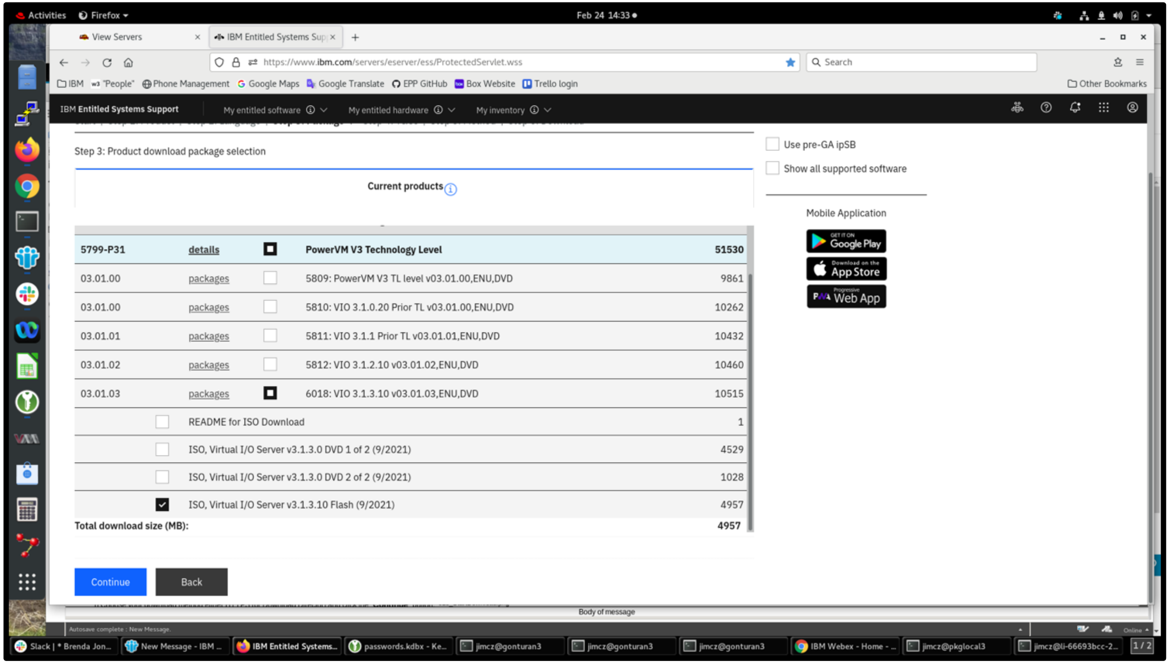

Step 3: Product download package selection

Current products

5799-P31	details	■	**PowerVM V3 Technology Level**	**51530**
03.01.00	packages	□	5809: PowerVM V3 TL level v03.01.00,ENU,DVD	9861
03.01.00	packages	□	5810: VIO 3.1.0.20 Prior TL v03.01.00,ENU,DVD	10262
03.01.01	packages	□	5811: VIO 3.1.1 Prior TL v03.01.01,ENU,DVD	10432
03.01.02	packages	□	5812: VIO 3.1.2.10 v03.01.02,ENU,DVD	10460
03.01.03	packages	■	6018: VIO 3.1.3.10 v03.01.03,ENU,DVD	10515
	□	README for ISO Download		1
	□	ISO, Virtual I/O Server v3.1.3.0 DVD 1 of 2 (9/2021)		4529
	□	ISO, Virtual I/O Server v3.1.3.0 DVD 2 of 2 (9/2021)		1028
	☑	ISO, Virtual I/O Server v3.1.3.10 Flash (9/2021)		4957

Total download size (MB): **4957**

□ Use pre-GA ipSB

□ Show all supported software

Mobile Application

Figure 14: 5799-P31 6018 Flash

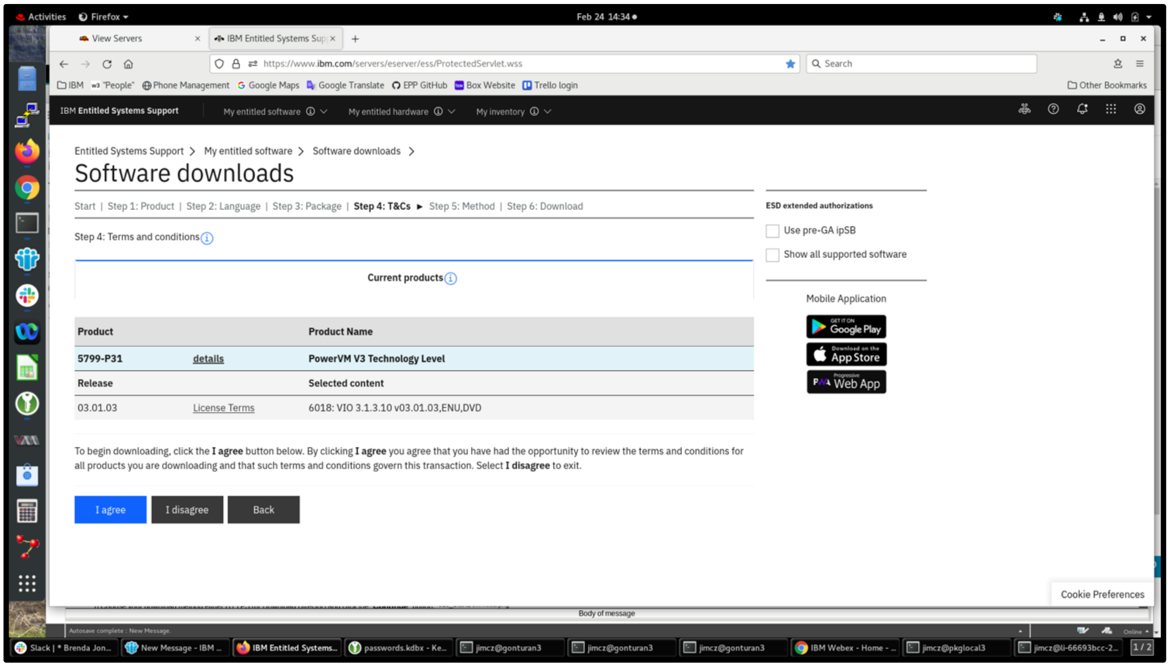

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start | Step 1: Product | Step 2: Language | Step 3: Package | **Step 4: T&Cs** ▸ Step 5: Method | Step 6: Download

Step 4: Terms and conditions

Current products

Product		Product Name
5799-P31	details	PowerVM V3 Technology Level
Release		**Selected content**
03.01.03	License Terms	6018: VIO 3.1.3.10 v03.01.03,ENU,DVD

To begin downloading, click the **I agree** button below. By clicking **I agree** you agree that you have had the opportunity to review the terms and conditions for all products you are downloading and that such terms and conditions govern this transaction. Select **I disagree** to exit.

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Figure 15: License Terms

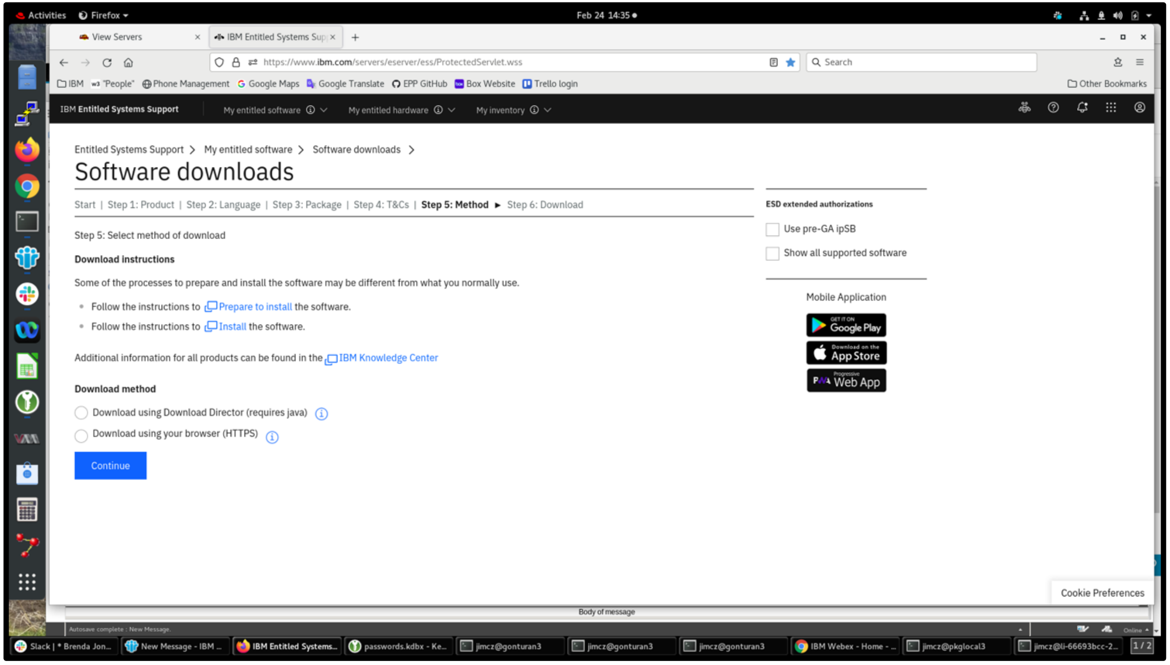

Entitled Systems Support > My entitled software > Software downloads >

Software downloads

Start | Step 1: Product | Step 2: Language | Step 3: Package | Step 4: T&Cs | **Step 5: Method** ▸ Step 6: Download

Step 5: Select method of download

Download instructions

Some of the processes to prepare and install the software may be different from what you normally use.

- Follow the instructions to Prepare to install the software.
- Follow the instructions to Install the software.

Additional information for all products can be found in the IBM Knowledge Center

Download method

○ Download using Download Director (requires java)

○ Download using your browser (HTTPS)

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Figure 16: Download Method

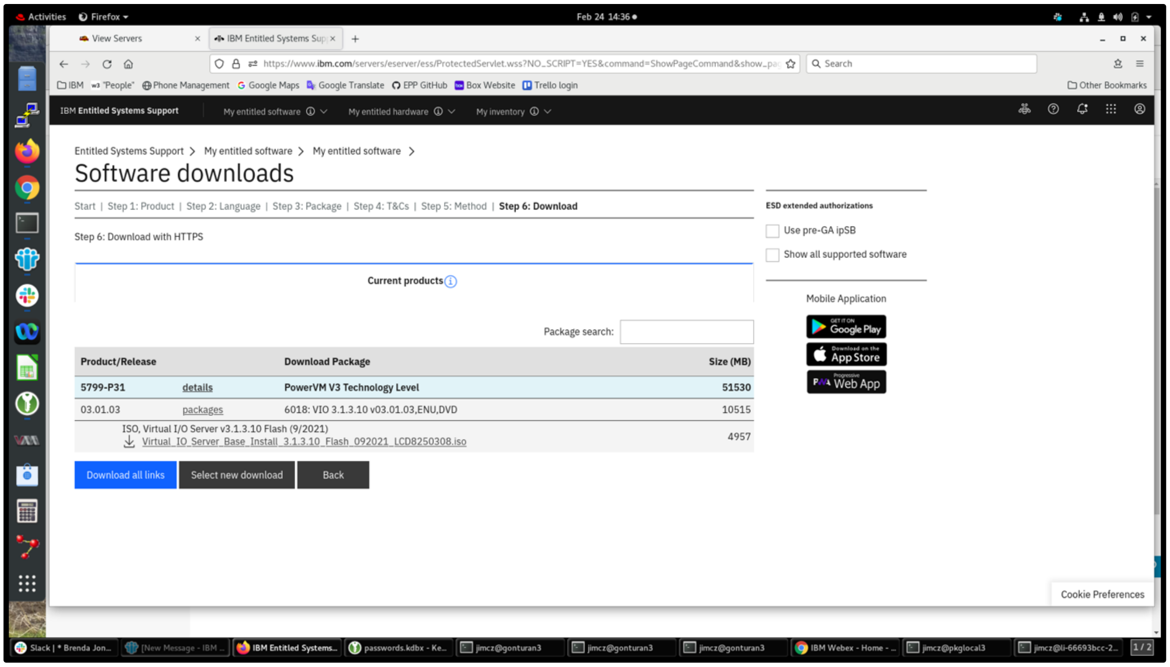

Entitled Systems Support > My entitled software > My entitled software >

Software downloads

Start | Step 1: Product | Step 2: Language | Step 3: Package | Step 4: T&Cs | Step 5: Method | **Step 6: Download**

Step 6: Download with HTTPS

Current products

Package search:

Product/Release		Download Package	Size (MB)
5799-P31	details	**PowerVM V3 Technology Level**	**51530**
03.01.03	packages	6018: VIO 3.1.3.10 v03.01.03,ENU,DVD	10515
		ISO, Virtual I/O Server v3.1.3.10 Flash (9/2021) Virtual_IO_Server_Base_Install_3.1.3.10_Flash_092021_LCD8250308.iso	4957

ESD extended authorizations

□ Use pre-GA ipSB

□ Show all supported software

Figure 17: Start Download

Power Systems

Virtual I/O Server

Note

Before using this information and the product it supports, read the information in "Notices" on page 297.

This edition applies to IBM® Virtual I/O Server Version 3.1.2, and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Virtual I/O Server

You can manage the Virtual I/O Server (VIOS) and client logical partitions by using the Hardware Management Console (HMC) and the Virtual I/O Server command-line interface.

The PowerVM® Editions feature includes the installation media for the VIOS software. The VIOS facilitates the sharing of physical I/O resources between client logical partitions within the server.

When you install the VIOS in a logical partition on a system that is managed by the HMC, you can use the HMC and the Virtual I/O Server command-line interface to manage the Virtual I/O Server and client logical partitions.

When you install the VIOS on a managed system and there is no HMC attached to the managed system when you install the VIOS, then the VIOS logical partition becomes the management partition. In POWER7® and POWER8® processor-based servers, the management partition provides the Integrated Virtualization Manager (IVM) web-based system management interface and a command-line interface that you can use to manage the system. IVM is not supported on POWER9™ processor-based servers.

For the most recent information about devices that are supported on the VIOS and to download VIOS fixes and updates, see the Fix Central website (http://www-933.ibm.com/support/fixcentral/).

Related information

PowerVM Information Roadmap
Virtual I/O Server commands

What's new in Virtual I/O Server

Read about new or changed information in Virtual I/O Server (VIOS) since the previous update of this topic collection.

September 2021

The following topics were updated:

- “Scenario: Configuring Shared Ethernet Adapter failover” on page 59
- “Scenario: Configuring Shared Ethernet Adapter failover without using a dedicated control channel adapter” on page 62
- “Shared Ethernet Adapter failover” on page 83
- “Shared Ethernet adapters for load sharing” on page 84

August 2021

The following topics were updated:

- “Limitations and restrictions for IBM i client logical partitions” on page 89

November 2020

The following information is a summary of the updates made to this topic collection:

- Added the topic “NPIV Multiple-Queue support” on page 9 with information about the NPIV multiple queue feature.
- Updated the topic “Disk” on page 21 with information about the virtual SCSI device read or write command timeout feature.
- Updated the topic “Networking considerations for shared storage pools” on page 124 with information about utilizing primary network interface and information about the limitations of using virtual IP address (VIPA).

April 2020

Replaced the information about the supported models with a reference to System software maps in the topic "Limitations and restrictions for IBM i client logical partitions" on page 89.

December 2019

The following topics were added or updated with information about networking considerations and restrictions for shared storage pools:

- "Configuring the system to create shared storage pools" on page 122
- "Networking considerations for shared storage pools" on page 124

October 2019

Added information about multiple Internet Small Computer Systems Interface (iSCSI) initiator support in the topic "iSCSI disk support for VIOS" on page 27.

July 2019

Added information about new attributes for Shared Ethernet Adapters (SEA) in the topic "Network attributes" on page 266.

August 2018

The following information is a summary of the updates made to this topic collection:

- Added information about the Internet Small Computer System Interface (iSCSI) disk support in VIOS in the topic "iSCSI disk support for VIOS" on page 27.
- Added information about the VIOS upgrade tool in the topic "Migrating the Virtual I/O Server by using the viosupgrade command or by using the manual method" on page 100.
- Added information about Shared Storage Pool (SSP) being migrated to the PostgreSQL database in the topic "Getting started with shared storage pools by using the VIOS command line interface" on page 121.
- Removed or updated obsolete information in various topics.
- Miscellaneous updates were made to this topic collection.

Virtual I/O Server overview

Learn the concepts of the Virtual I/O Server (VIOS) and its primary components.

The VIOS is part of the PowerVM Editions hardware feature. The VIOS is a software that is located in a logical partition. This software facilitates the sharing of physical I/O resources between client logical partitions within the server. The VIOS provides virtual Small Computer Serial Interface (SCSI) target, virtual Fibre Channel, Shared Ethernet Adapter, and PowerVM Active Memory Sharing capability to client logical partitions within the system. The VIOS also provides the Suspend/Resume feature to AIX®, IBM i, and Linux® client logical partitions within the system when you are managing a POWER7, POWER8, or POWER9 processor-based server.

Note: The Suspend/Resume feature of logical partitions is not supported on the POWER9 Power Systems servers. This feature is supported on other models of Power Systems servers, with appropriate levels of the management console, firmware, and PowerVM.

As a result, you can perform the following functions on client logical partitions:

- Share SCSI devices, Fibre Channel adapters, Ethernet adapters

- Expand the amount of memory available to logical partitions and suspend and resume logical partition operations by using paging space devices when you are managing a POWER7, POWER8, or POWER9 processor-based server.

A dedicated logical partition is required for the VIOS software solely for its use.

You can use the VIOS to perform the following functions:

- Sharing of physical resources between logical partitions on the system
- Creating logical partitions without requiring additional physical I/O resources
- Creating more logical partitions than there are I/O slots or physical devices available with the ability for logical partitions to have dedicated I/O, virtual I/O, or both
- Maximizing use of physical resources on the system
- Helping to reduce the storage area network (SAN) infrastructure

Related information

Virtual I/O Server commands

Operating system support for VIOS client logical partitions

For more information about operating systems that run on client logical partitions and that are supported by the Virtual I/O Server (VIOS), see System software maps.

Virtual Fibre Channel

With *N_Port ID Virtualization (NPIV)*, you can configure the managed system so that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter.

To access physical storage in a typical storage area network (SAN) that uses Fibre Channel, the physical storage is mapped to logical units (LUNs) and the LUNs are mapped to the ports of physical Fibre Channel adapters. Each physical port on each physical Fibre Channel adapter is identified using one worldwide port name (WWPN).

NPIV is a standard technology for Fibre Channel networks that enables you to connect multiple logical partitions to one physical port of a physical Fibre Channel adapter. Each logical partition is identified by a unique WWPN, which means that you can connect each logical partition to independent physical storage on a SAN.

To enable NPIV on the managed system, you must complete the following steps:

- Create a Virtual I/O Server logical partition (version 2.1, or later) that provides virtual resources to client logical partitions.
- Assign the physical Fibre Channel adapters (that support NPIV) to the Virtual I/O Server logical partition.
- Connect virtual Fibre Channel adapters on the client logical partitions to virtual Fibre Channel adapters on the Virtual I/O Server logical partition.

A *virtual Fibre Channel adapter* is a virtual adapter that provides client logical partitions with a Fibre Channel connection to a storage area network through the Virtual I/O Server logical partition. The Virtual I/O Server logical partition provides the connection between the virtual Fibre Channel adapters on the Virtual I/O Server logical partition and the physical Fibre Channel adapters on the managed system.

The following figure shows a managed system configured to use NPIV.

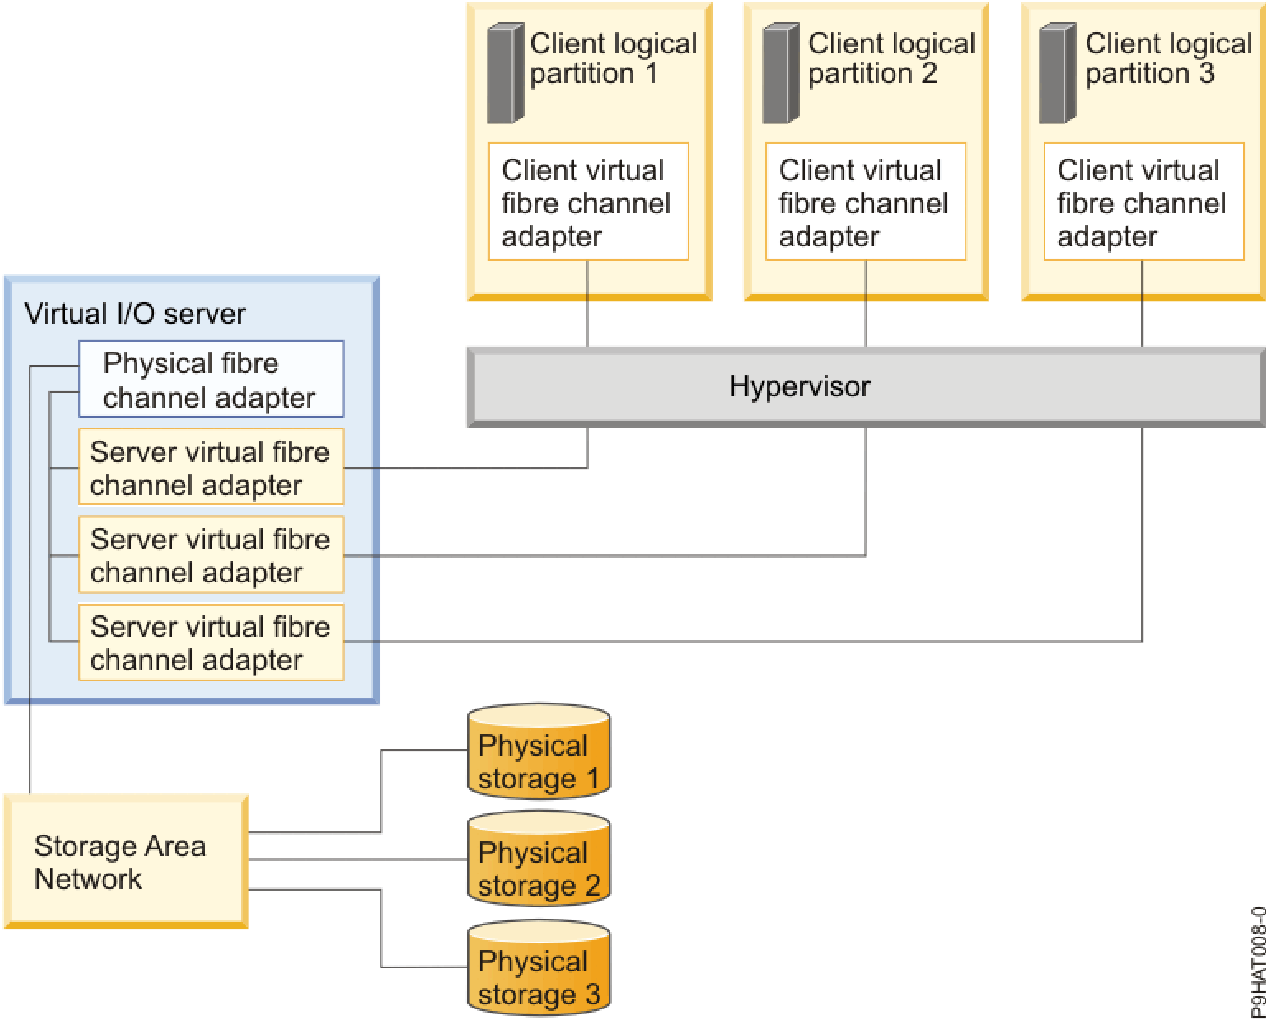
Client logical partition 1
Client virtual fibre channel adapter
Client logical partition 2
Client virtual fibre channel adapter
Client logical partition 3
Client virtual fibre channel adapter
Virtual I/O server
Physical fibre channel adapter
Server virtual fibre channel adapter
Server virtual fibre channel adapter
Server virtual fibre channel adapter
Hypervisor
Storage Area Network
Physical storage 1
Physical storage 2
Physical storage 3
P9HAT008-0

The figure shows the following connections:

- A storage area network (SAN) connects three units of physical storage to a physical Fibre Channel adapter that is located on the managed system. The physical Fibre Channel adapter is assigned to the Virtual I/O Server and supports NPIV.
- The physical Fibre Channel adapter connects to three virtual Fibre Channel adapters on the Virtual I/O Server. All three virtual Fibre Channel adapters on the Virtual I/O Server connect to the same physical port on the physical Fibre Channel adapter.
- Each virtual Fibre Channel adapter on the Virtual I/O Server connects to one virtual Fibre Channel adapter on a client logical partition. Each virtual Fibre Channel adapter on each client logical partition receives a pair of unique WWPNs. The client logical partition uses one WWPN to log into the SAN at any given time. The other WWPN is used when you move the client logical partition to another managed system.
- In this case, Client logical partition 1 accesses Physical storage 1, Client logical partition 2 accesses Physical storage 2, and Client logical partition 3 accesses Physical storage 3.

For IBM® i client partitions, the LUNs of the physical storage connected with NPIV require a storage-specific device driver and do not use the generic virtual SCSI device driver. The Virtual I/O Server cannot access and does not emulate the physical storage to which the client logical partitions have access. The Virtual I/O Server provides the client logical partitions with a connection to the physical Fibre Channel adapters on the managed system.

Note: The Virtual I/O Server cannot access and does not emulate the physical storage to which the client logical partitions have access.

There is always a one-to-one relationship between virtual Fibre Channel adapters on the client logical partitions and the virtual Fibre Channel adapters on the Virtual I/O Server logical partition. That is, each virtual Fibre Channel adapter on a client logical partition must connect to only one virtual Fibre Channel adapter on the Virtual I/O Server logical partition, and each virtual Fibre Channel on the Virtual I/O Server logical partition must connect to only one virtual Fibre Channel adapter on a client logical partition.

Note: Mapping of multiple Virtual Fibre Channel adapters of a single client logical partition through multiple virtual server Fibre Channel adapters to the same physical Fibre Channel adapter is not recommended.

Using SAN tools, you can zone and mask LUNs that include WWPNs that are assigned to virtual Fibre Channel adapters on client logical partitions. The SAN uses WWPNs that are assigned to virtual Fibre Channel adapters on client logical partitions the same way it uses WWPNs that are assigned to physical ports.

The following operating system (OS) levels are supported for client logical partitions to configure VFC adapters.

Table 1. Allowed OS levels for client logical partitions to configure VFC adapters

Operating system	**Supported versions**
AIX®	Version 5.3 Technology Level 9 Version 6.1 Technology Level 2, or later
IBM® i	Version 6.1.1, or later
SUSE Linux Enterprise Server	Version 10 service pack 3, or later Version 11, or later
Red Hat Enterprise Server	Version 5.4, or later Version 6, or later

Virtual Fibre Channel for HMC-managed systems

On systems that are managed by the Hardware Management Console (HMC), you can dynamically add and remove virtual Fibre Channel adapters to and from the Virtual I/O Server logical partition and each client logical partition. You can also view information about the virtual and physical Fibre Channel adapters and the worldwide port names (WWPNs) by using Virtual I/O Server commands.

To enable N_Port ID Virtualization (NPIV) on the managed system, you create the required virtual Fibre Channel adapters and connections as follows:

- You use the HMC to create virtual Fibre Channel adapters on the Virtual I/O Server logical partition and associate them with virtual Fibre Channel adapters on the client logical partitions.
- You use the HMC to create virtual Fibre Channel adapters on each client logical partition and associate them with virtual Fibre Channel adapters on the Virtual I/O Server logical partition. When you create a virtual Fibre Channel adapter on a client logical partition, the HMC generates a pair of unique WWPNs for the client virtual Fibre Channel adapter.
- You can connect the virtual Fibre Channel adapters on the Virtual I/O Server to the physical ports of the physical Fibre Channel adapter by running the **`vfcmap`** command on the Virtual I/O Server.

The HMC generates WWPNs based on the range of names available for use with the prefix in the vital product data on the managed system. This 6–digit prefix comes with the purchase of the managed system and includes 32,000 pairs of WWPNs. When you remove a virtual Fibre Channel adapter from a client logical partition, the hypervisor deletes the WWPNs that are assigned to the virtual Fibre Channel adapter on the client logical partition. The HMC does not reuse the deleted WWPNs when generating WWPNs for virtual Fibre Channel adapters in the future. If you run out of WWPNs, you must obtain an activation code that includes another prefix with another 32,000 pairs of WWPNs.

To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client logical partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client logical partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.

You can dynamically add and remove virtual Fibre Channel adapters to and from the Virtual I/O Server logical partition and to and from client logical partitions.

Table 2. Dynamic partitioning tasks and results for virtual Fibre Channel adapters

Dynamically add or remove virtual Fibre Channel adapter	**To or from a client logical partition or a Virtual I/O Server logical partition**	**Result**
Add a virtual Fibre Channel adapter	To a client logical partition	The HMC generates the a pair of unique WWPNs for the client virtual Fibre Channel adapter.
Add a virtual Fibre Channel adapter	To a Virtual I/O Server logical partition	You need to connect the virtual Fibre Channel adapter to a physical port on a physical Fibre Channel adapter.
Remove a virtual Fibre Channel adapter	From a client logical partition	• The hypervisor deletes the WWPNs and does not reuse them. • You must either remove the associated virtual Fibre Channel adapter from the Virtual I/O Server, or associate it with another virtual Fibre Channel adapter on a client logical partition.
Remove a virtual Fibre Channel adapter	From a Virtual I/O Server logical partition	• The Virtual I/O Server removes the connection to the physical port on the physical Fibre Channel adapter. • You must either remove the associated virtual Fibre Channel adapter from the client logical partition, or associate it with another virtual Fibre Channel adapter on the Virtual I/O Server logical partition.

The following table lists the Virtual I/O Server commands that you can run to view information about the Fibre Channel adapters.

Table 3. Virtual I/O Server commands that display information about Fibre Channel adapters

Virtual I/O Server command	Information displayed by command
lsmap	• Displays the virtual Fibre Channel adapters on the Virtual I/O Server that are connected to the physical Fibre Channel adapter • Displays attributes of the virtual Fibre Channel adapters on the client logical partitions that are associated with the virtual Fibre Channel adapters on the Virtual I/O Server that are connected to the physical Fibre Channel adapter
lsnports	Displays information about the physical ports on the physical Fibre Channel adapters that support NPIV, such as: • The name and location code of the physical port • The number of available physical ports • The total number of WWPNs that the physical port can support • Whether the switches, to which the physical Fibre Channel adapters are cabled, support NPIV

You can also run the **lshwres** command on the HMC to display the remaining number of WWPNs and to display the prefix that is currently used to generate the WWPNs.

NPIV disk validation for Live Partition Migration

This topic provides information about the logical unit (LU) level validation for migration of N_Port ID Virtualization (NPIV) clients. During the validation phase of Live Partition Migration (LPM), checks are performed to ensure that the NPIV client has access to the same set of LUs on both the destination server and the source server. These checks can be optionally enabled on source and destination Virtual I/O Server (VIOS). Only block storage devices are checked for compatibility and other devices are skipped.

Disk validation can add considerable time to N_Port ID Virtualization (NPIV) mobility. The time spent depends on the number of devices you have mapped to a client partition. The time spent might impact the maintenance windows and you might want to consider validating the NPIV disk periodically, perform disk validation tasks outside the maintenance windows or just before a maintenance window.

Disk validation might fail if your storage area network (SAN) is more unstable than the earlier versions of VIOS in which a VIOS, only validated access to target ports. This is because more commands are sent through the SAN to devices.

New attributes are added to the `vioslpm0` device of the VIOS to enable or disable LU level validation. The source and destination VIOS must both support disk mapping validation regardless of the *src_lun_val* attribute for NPIV disk validation to find configuration errors. If a source VIOS generates the appropriate data stream and the destination VIOS is not capable of disk validation, the additional disk information is ignored by the destination VIOS. Consider this scenario while scheduling VIOS maintenance.

NPIV disk validation is not supported on HMC Version 7 Release 7.4.4, or earlier. Timer values used in these versions of HMC might cause validation issues. Consider this restriction before enabling disk validation.

Use of src_lun_val in the HMC

Disk mapping validation is performed only during validation; it is not performed during migration. In the migration phase, only port validation is performed. If you are using the HMC graphical user interface, you must perform validation for each LPM operation. Consider this restriction before enabling disk validation by changing the *src_lun_val* attribute, particularly if you are using an inordinate number of disks and if you are using the HMC.

If you are using the HMC migration command, validation is performed only if the **–o** flag is set to the character *v* and migration is performed only if the –*o* flag is set to the character *m*. They are mutually exclusive.

You can choose to use the HMC command line to control when validation occurs in relation to maintenance windows and always enable disk validation on the VIOS. This feature is useful if you are already performing validation from the command line and want to perform disk mapping validation for users with very large configurations, for example a user with 4,000 to 5,000 disks.

Attributes for NPIV disk validation

The following attributes can be used during NPIV disk validation.

Table 4. Attributes for NPIV disk validation

Attribute name	Description
src_lun_val	This attribute can be set to *off* or *on* by using the **chdev** command. The default value is *off* so that behavior is not changed during the NPIV LPM validation. This means that if the value is set to *off*, disk mapping is not validated. To turn on disk mapping validation, run the following command: `chdev -dev vioslpm0 -attr src_lun_val=on`
dest_lun_val	This attribute can be changed to several different values by using the **chdev** command. The default value is *restart_off*. The attribute can be set to the following values: **restart_off** If this attribute is set to *restart_off*, disk mapping LPM validation depends on the data stream generated by the source VIOS. Disk mapping validation is not performed for suspend and resume operations, regardless of the source data stream. Use this attribute value when the data streams stored for a particular client are more likely to be stale than data streams collected at the time of LPM validation.
	lpm_off If this attribute is set to *lpm_off* , disk mapping LPM validation is turned off, regardless of the data stream generated by the source. VIOS. Disk mapping validation performed for suspend resume operations depends on the source VIOS data stream.
	on If this attribute is set to *on*, disk mapping validation completely depends on the data stream generated by the source VIOS.

Table 4. Attributes for NPIV disk validation (continued)

Attribute name	Description
	off If this attribute is set to *off*, disk mapping validation is not performed for any operation.
max_val_cmds	This attribute allows you to change the number of commands that are allocated for NPIV disk validation. The commands are used to discover the identity of each disk that the client can access. Threads are allocated groups of work and the group size depends on available commands. If more work is completed, validation completes sooner. Commands require VIOS memory resource. If more commands are allocated, more bandwidth is used per physical port on the destination VIOS. From the physical port, a particular virtual NPIV server adapter is used to access the SAN on behalf of the client. You might not need to change this value, unless you have an aberrant configuration.

NPIV Multiple-Queue support

Learn about the modernization of *N_Port ID Virtualization (NPIV)* by enabling multiple-queues, which is commonly known as NPIV Multiple-Queue (MQ).

Currently, Fibre Channel (FC) adapters with high bandwidth, such as 16 GB or 32 GB FC adapters support multiple-queue pairs for storage I/O communication. Multiple-queue pairs in the physical FC stack significantly improve the input/output requests per second (IOPS) due to the ability to drive the I/Os in parallel through the FC adapter. The objective of the NPIV Multiple-Queue is to add similar Multiple-Queue support to all components such as the client operating system (OS), POWER® Hypervisor (PHYP), and the Virtual I/O Server (VIOS). The NPIV VIOS stack and the PHYP are updated to allow client LPARs to access multiple-queues. The Multiple-Queue feature is supported only on AIX client logical partitions and on VIOS Version 3.1.2, or later.

NPIV scaling improvements through Multiple-Queue provides the following benefits:

- Efficient utilization of available Multiple-Queue FC adapters bandwidth when mapped to a single or multiple LPARs.
- Enable and drive multiple logical units (LUN) level I/O traffic in parallel through FC adapter queues.
- Storage I/O performance improvement due to increased input/output requests per second (IOPS).

The following figure shows a managed system that is configured to use NPIV Multiple-Queues:

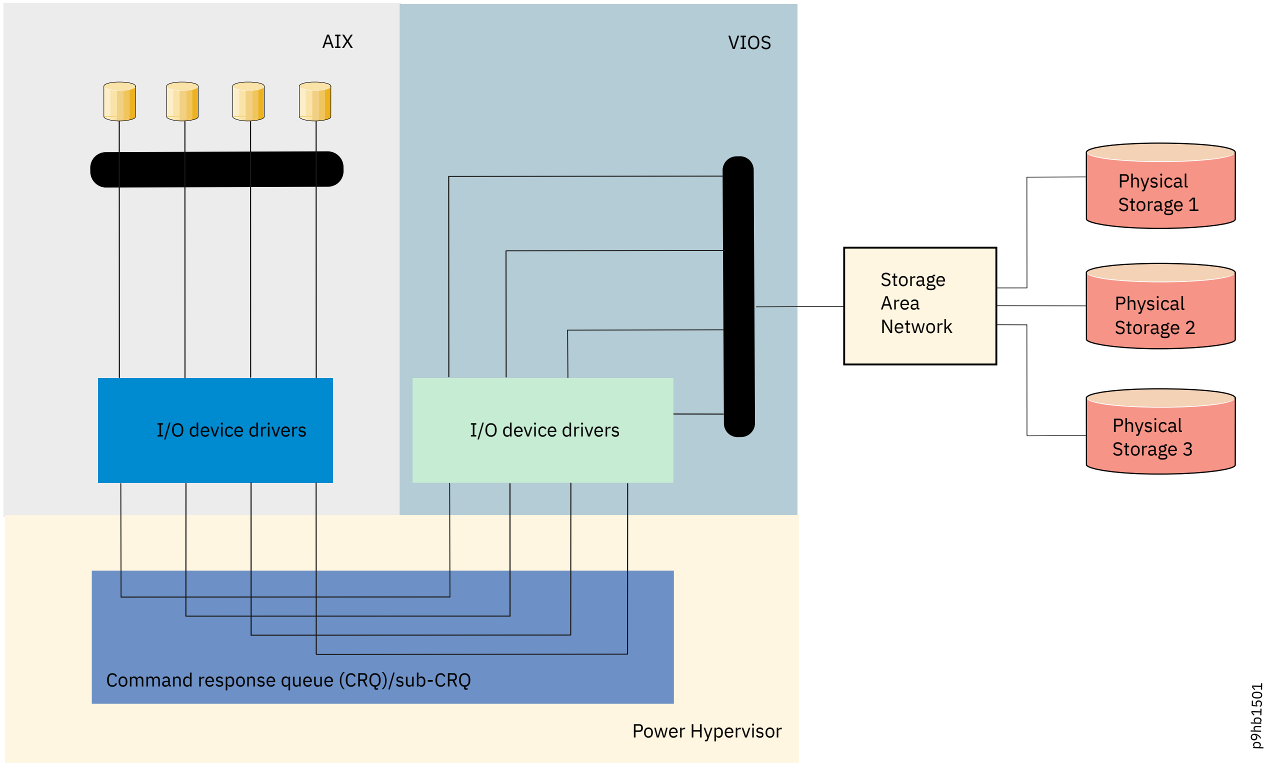

Hardware support and requirements to enable the Multiple-Queue feature for NPIV

Table 5. Multiple-queue for NPIV

Operating system/PFW	**Supported versions**
Hardware	POWER9 processor-based systems
AIX	Version 7.2 Technology Level 05, or later
VIOS	Version 3.1.2, or later
POWER firmware	Version 940, or later
Fibre Channel (FC) adapter	Emulex FC 16 or 32 Gb FC adapters or any high-bandwidth Fibre Channel adapters that support Multiple-Queue feature.
IBM i	Not supported
Linux Enterprise Server (SUSE, Red Hat®)	Not supported

Performance benefits

NPIV Multiple-Queue enablement provides improved storage I/O performance for different types of workloads.

LPAR mobility in a Multiple-Queue supported environment

NPIV Multiple-Queue enablement for all components requires support from the client operating system, hypervisor, and VIOS. During the LPM operation, if either hypervisor or the VIOS of the destination system does not support Multiple-Queue, Multiple-Queue is not enabled after the LPM operation.

LPAR mobility in a Multiple-Queue supported environment section is described based on the following perspectives:

- The Multiple-Queue feature is supported only on AIX client logical partitions and on VIOS Version 3.1.2, or later.
- LPM from a VIOS perspective, considering the potential implementations of other PowerVM clients.
- LPM and Multiple-Queue from a firmware perspective.
- **Considerations for NPIV configuration and LPM validation**
  - During the initial configuration, when you connect the NPIV client to the VIOS, the VIOS reports whether the Multiple-Queue feature is supported. If the feature is supported, VIOS reports whether it can migrate from an environment where it has established multiple queues to the destination where fewer queues can be established. The VIOS also reports whether it can continue to perform I/O operations in a single-queue environment (systems with VIOS version earlier than 3.1.2).
  - Power® firmware supports the Multiple-Queue feature through the implementation of a construct called Subordinate Command Response Queues (sub-CRQs). The NPIV sub-CRQ construct is supported on POWER9, or later systems. The sub-CRQ construct is lost if a client is moved from a POWER9 system to an earlier model POWER system, or if a client is moved to systems with older firmware levels than the current system.
  - During the initial configuration, the VIOS provides information about the firmware and adapters so that the NPIV client can determine whether to maintain NPIV sub-CRQ construct that support the Multiple-Queue feature. During the LPM operation, if the firmware moves the sub-CRQ construct from the source managed system to the destination managed system, the NPIV client can store the queue resources and use it later when the LPM operation is performed on an environment where all the resources are available.
- **LPM scenarios and Multiple-Queue behavior in an AIX client**
  - During the initial configuration of NPIV client, the AIX NPIV client LPAR exchanges capabilities with the VFC host such as Multiple-Queue, migration, and firmware levels and then performs the configuration. These capabilities are exchanged again during the LPM operation at the destination managed system. The Multiple-Queue feature is enabled or deprecated based on these capabilities.
  - When the AIX LPAR is migrated from the source system with the NPIV Multiple-Queue support setup to the destination system with NPIV Multiple-Queue support setup, the NPIV stack continues to run in the Multiple-Queue environment:
    - The performance might remain the same until the NPIV client can create the same number of queues and has similar FC adapter bandwidth that is available at the destination system when compared to the source managed system.
    - While exchanging the initial capabilities during the LPM operation, if the VFC host at the destination managed system reports less queues as compared to the number of queues that are configured on the source managed system, the NPIV client configures and continues sending I/O requests through these available queues.
    - The performance might be impacted if either of queues on the source managed system or the destination managed system is less, or if the storage bandwidth at the destination managed system is less when compared to the source managed system.
    - While exchanging the initial capabilities during the LPM operation, if the VFC host at the destination managed system reports more queues, the NPIV client uses the same number of queues when compared to the number of queues that are configured on the source managed system.

      Examples:

      - If the number of queues that are configured at the source managed system is 8 and if the VFC host at the destination managed system reports 4 queues, only 4 queues are configured at the destination managed system. If the same LPAR is migrated back or to another destination system where the VFC host reports 8 queues, the NPIV client is configured with 8 queues.
      - If the number of queues that are configured at the source managed system is 8 and if the VFC host at the destination managed system reports 16 queues, the VFC client continues to run with 8 queues.

- When the AIX LPAR is migrated from the NPIV Multiple-Queue environment to a managed system with an older firmware level, the Multiple-Queue resources are lost and the performance might reduce regardless of the adapters in the destination managed system. The NPIV client does not establish Multiple-Queue when it is subsequently moved to a system that supports the Multiple-Queue environment.
- When the AIX LPAR is migrated to an environment where the VIOS and the firmware support Multiple-Queue, but the FC adapters, such as 4 or 8 Gb Emulex, do not support Multiple-Queue, the sub-queue resources are retained by the AIX client. The AIX client can be used if the client is subsequently moved to an environment that supports Multiple-Queue. Performance issues might occur after migrating from a Multiple-Queue environment to an environment that does not support Multiple-Queue.
- When the AIX LPAR is migrated to an environment where the VIOS is not capable of the Multiple-Queue feature, the sub-queues are lost and multiple queues are deprecated. The NPIV client runs in a single queue mode (similar to the NPIV setup in AIX 7200-04 or earlier, and VIOS Version 3.1.1, or earlier versions). The NPIV client does not establish multiple queues when it is subsequently moved to a system that supports Multiple-Queue environment.
- When a Multiple-Queue NPIV client partition (AIX 7200-05, or later) is migrated from a POWER8 or POWER7 system to a POWER9 system with Multiple-Queue setup, the partition continues to operate in the NPIV single-channel mode because after you migrate a partition from a lower processor compatibility mode to a POWER9 system, the partition continues to run in a lower processor compatibility mode of POWER8 or POWER7 systems. When the partition is booted with native mode on a POWER9 system, the NPIV Multiple-Queue is enabled during the NPIV configuration as part of the startup process.

  Note: POWER firmware level FW930, or later supports the sub-CRQ construct that is used for Multiple-Queue enablement. Hence, performing the LPM operation from a Multiple-Queue aware setup to a system with POWER firmware level FW930, or later and VIOS Version 3.1.2, or later preserves the sub-CRQ construct. Migrating this LPAR back to Multiple-Queue aware setup enables the Multiple-Queue feature.

VIOS Tunable Attributes

New VIOS tunable attributes are available in VIOS version 3.1.2, or later as part of the NPIV Multiple-Queue feature to provide flexibility with the number of FC adapter queues (physical queues) that each VFC host adapter uses. The NPIV Multiple-Queue feature also provides QoS type features and tunable attributes that are applicable to all the VFC host adapters. The ***num_per_range*** attribute can be set at the VIOS partition level and can be overridden at the individual VFC host adapter level.

The NPIV Multiple-Queue supports a new pseudo device called **viosnpiv0**. The partition wide tunable attributes are provided by the **viosnpiv0** device. The local tunable attributes are provided by the VFC host adapter device. The following tables describe various tunable attributes that can be used for optimal performance:

Table 6. viosnpiv0 device attributes

Attribute	Min value	Max value	Default value	Description
num_per_range	4	64	8	A VIOS level tunable attribute. It indicates the number of FC SCSI queues that each VFC host uses.
num_local_cmds	1	64	5	Allows you to trade off memory resources and performance. A higher value might improve performance for fewer I/O workloads. It controls resources that are allocated for each specific queue that is in use by the VFC host adapter.

Table 6. viosnpiv0 device attributes (continued)

Attribute	Min value	Max value	Default value	Description
bufs_per_cmd	1	64	10	Allows you to trade off memory resource and performance. A higher value might improve performance for larger I/O workloads.

Table 7. vfchost attributes

Attribute name	Min value	Max value	Default value	Description
num_per_range	4	64	0	If this attribute is set to a nonzero value, it overrides the partition wide *num_per_range* attribute of the **viosnpiv0** device. If the attribute value is 0, this tunable attribute is not in effect.
limit_intr	Boolean (true or false)	Boolean (true or false)	false	A local tunable attribute. If this attribute is set to *true*, it is expected to negatively impact the performance for a particular adapter. It reduces the number of processors and IOPS that are used to service the VFC host adapter. It takes precedence over the *num_per_range* attribute.
label	N/A	N/A	""	Used to tag a VFC host adapter with a user-defined string identifier. After a successful LPM operation, the VFC host adapter on the destination VIOS will have the same label as the source VIOS.

Note: The attributes that are related to Multiple-Queue are lost if you are moving from a VIOS that supports Multiple-Queue to another VIOS that does not support Multiple-Queue (if the NPIV client is capable of such a mobility operation).

The local ***limit_intr*** attribute has the highest precedence. If the ***limit_intr*** is set to *false*, the local attribute ***num_per_range*** is effective. When the local ***num_per_range*** attribute is not set, the partition wide attribute ***num_per_range*** is effective.

The number of queues that a client uses depends on the FC adapter, FW level, and client capabilities, and also on the VIOS level and the tunable attributes of the VFC host adapter. After a successful LPM operation, if the client is using multiple queues, the local attribute ***num_per_range*** or the ***limit_intr*** attribute of the VFC host adapter is set on the destination managed system that is based on the value that is used at the source managed system.

Table 8. AIX VFC client tunable attributes

Attribute name	Min value	Max value	Default value	Description
lg_term_dma	1 MB	16 MB	8 MB	Indicates the memory that is required by the virtual driver for its internal data structure. This attribute value can be modified or increased for the environment with large number of NPIV disks.

Table 8. AIX VFC client tunable attributes (continued)

Attribute name	Min value	Max value	Default value	Description
max_xfer_size	1 MB	16 MB	1 MB	Allows you to set the maximum transfer size for single I/O. This tunable attribute must be modified to suit I/O transfer size in different environments. For example, Tape drives (sequential I/O) use large block sizes for I/O transfers.
num_cmd_elems	20	2048	1024	Determines the maximum number of active I/O operations at any given point of time.
num_io_queues	1	16	8	Determines the number of I/O queues that are used in the SCSI I/O communication.
label	N/A	N/A	""	User-defined name to identify the adapter.
num_sp_cmd_elem	512	2048	512	Determines the maximum number of special command operations at any given point of time.

Notes:

- The number of queues that the NPIV client uses depends on several factors such as FC adapter, FW level, VIOS level, and tunable attributes of the VFC host adapter. During the initial configuration, the VFC client negotiates the number of queues with the VFC host and configures the minimum value of ***num_io_queues*** attribute and the number of queues that are reported by the VFC host.
- After the initial configuration, the negotiated number is the maximum number of channels that the VFC client can enable. If the VFC host renegotiates more channels after operations (such as remap, VIOS restart, and so on), the number of channels remains the same as the initially negotiated number. However, if the VFC host renegotiates with fewer channels, the VFC client reduces its configured channels to this new lower number.

  For example, if the initial negotiated number of channels between the VFC client and VFC host is 8, and later if the VFC host renegotiates the number of channels as 16, the VFC client continues to run with 8 channels. If the VFC host renegotiates the number of channels as 4 channels, the VFC client adjusts its number of configured channels to 4. However, if the VFC host renegotiates the number of channels as 8 channels, which result in increasing the number of configured channels to 8, the VFC client must be reconfigured to renegotiate the number of channels from the client side.

Virtual SCSI

Using virtual Small Computer Serial Interface (SCSI), client logical partitions can share disk storage and tape or optical devices that are assigned to the Virtual I/O Server (VIOS) logical partition.

Physical storage devices such as disk, tape, Universal Serial Bus (USB) mass storage, or optical devices that are attached to the VIOS logical partition can be shared by one or more client logical partitions. The VIOS is a standard storage subsystem that provides standard logical unit numbers (LUNs) that are compliant with the SCSI. The VIOS can export a pool of heterogeneous physical storage as a homogeneous pool of block storage in the form of SCSI disks. The VIOS is a storage subsystem. Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices that are exported by the VIOS are limited to the domain within the server. Therefore, although the SCSI LUNs are SCSI-compliant, they might not meet the needs of all applications, particularly those applications that exist in a distributed environment.

The following SCSI peripheral device types are supported:

- Disk that is backed by logical volume
- Disk that is backed by physical volume
- Disk that is backed by file
- Disk that is backed by a logical unit in shared storage pools
- Optical CD-ROM, DVD-RAM, and DVD-ROM
- Optical DVD-RAM backed by file
- Tape devices
- USB mass storage devices

Virtual SCSI is based on a client-server relationship model as described in the following points.

- The VIOS owns the physical resources and the *virtual SCSI server adapter*, and acts as a server, or SCSI target device. The client logical partitions have a SCSI initiator referred to as the *virtual SCSI client adapter*, and accesses the virtual SCSI targets as standard SCSI LUNs.
- The configuration and provisioning of virtual disk resources can be performed by using the HMC or the VIOS command line.
- Physical disks owned by the VIOS can be exported and assigned to a client logical partition as a whole, added to a shared storage pool, or can be partitioned into parts, such as logical volumes or files. The logical volumes and files can then be assigned to different logical partitions. Therefore, by using virtual SCSI, you can share adapters and disk devices.
- Logical units in logical volumes and file-backed virtual devices prevent the client partition from participating in Live Partition Mobility. To make a physical volume, logical volume, or file available to a client logical partition requires that it must be assigned to a virtual SCSI server adapter on the VIOS. The client logical partition accesses its assigned disks through a virtual SCSI client adapter. The virtual SCSI client adapter recognizes standard SCSI devices and LUNs through this virtual adapter.

  Note: Logical units in logical volumes and file-backed virtual devices might prevent the client partition from participating in Live Partition Mobility.

Thin provisioning

Thin provisioning is applicable to logical units on Shared Storage Pools (SSP). On the VIOS, for logical units in shared storage pools, you can thin-provision a client virtual SCSI device for better storage space utilization. In a thin-provisioned device, the used storage space might be greater than the actual used storage space. If the blocks of storage space in a thin-provisioned device are unused, the device is not entirely backed by physical storage space. With thin-provisioning, the storage capacity of the storage pool can be exceeded. When the storage capacity is exceeded, a threshold exceeded alert is raised. To identify that a threshold alert has occurred, check the errors listed in the HMC serviceable events or the VIOS system error log by running the **`errlog`** command in the VIOS command line. To recover after the threshold has exceeded, you can add physical volumes to the storage pool. You can verify that the threshold is no longer exceeded in the HMC serviceable events or the VIOS system error log. For instructions on how to add physical volumes to the storage pool by using the VIOS command-line interface, see Adding physical volumes to the storage pool by using the VIOS command-line interface. For instructions on how to add physical volumes to the storage pool by using the VIOS configuration menu, see Adding physical volumes to the storage pool by using the VIOS configuration menu. You can also increase the storage capacity of the storage pool by deleting data.

Persistent reserve

On the VIOS, multiple applications running on the virtual client can manage reservations on virtual disks of the client by using the Persistent Reserves standard. These reservations persist across hard resets, logical unit resets, or initiator target nexus loss. Persistent reservations that are supported by logical devices from the VIOS shared storage pools support the required features for the SCSI-3 Persistent Reserves standard.

Thick provisioning

On the VIOS, you can thick-provision a virtual disk. In a thick-provisioned virtual disk, you can allocate or reserve storage space while initially provisioning the virtual disk. The allocated storage space for the thick-provisioned virtual disk is assured. This operation ensures that there are no failures because of lack of storage space. By using thick-provisioning, virtual disks have faster initial access time because the storage is already allocated.

The following figure shows a standard virtual SCSI configuration.

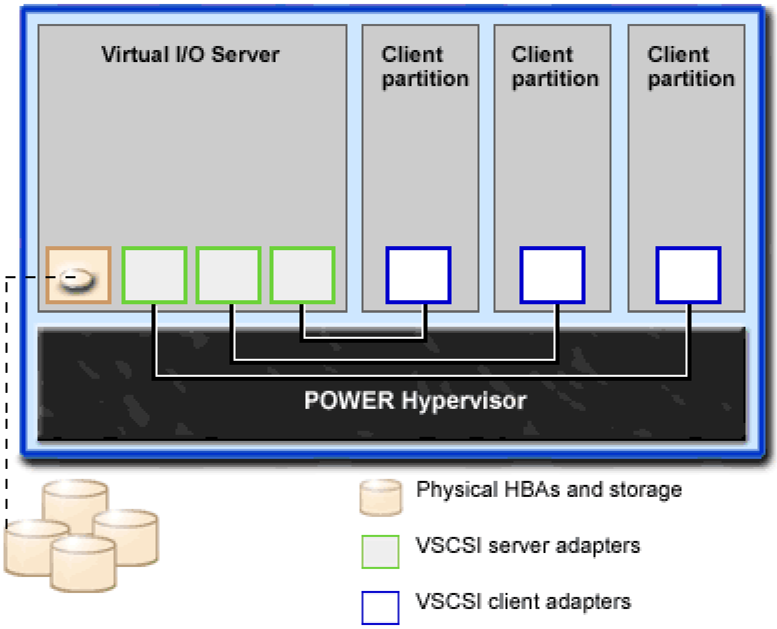

Note: The VIOS must be fully operational for the client logical partitions to be able to access virtual devices.

Related tasks

Adding physical volumes to the storage pool
You can add physical volumes to the storage pool by using the Virtual I/O Server (VIOS) command-line interface.

Virtual I/O Server storage subsystem overview

Learn about the Virtual I/O Server storage subsystem.

The Virtual I/O Server storage subsystem is a standard storage subsystem that provides standard logical unit numbers (LUNs) compliant with the Small Computer Serial Interface (SCSI). The Virtual I/O Server is a storage subsystem. Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices that are exported by the Virtual I/O Server are limited to the domain within the server.

Like typical disk storage subsystems, the Virtual I/O Server has a distinct front end and backend. The front end is the interface to which client logical partitions attach to view standard SCSI-compliant LUNs. Devices on the front end are called *virtual SCSI devices*. The backend is made up of physical storage resources. These physical resources include physical disk storage, both SAN devices and internal storage devices, optical devices, tape devices, logical volumes, and files.

To create a virtual device, some physical storage must be allocated and assigned to a virtual SCSI server adapter. This process creates a virtual device instance (vtscsi*X* or vtopt*X*). The device instance can be considered a mapping device. It is not a real device, but rather a mechanism for managing the mapping of the portion of physical backend storage to the front-end virtual SCSI device. This mapping device re-creates the physical-to-virtual allocations in a persistent manner when the Virtual I/O Server is restarted.

Physical storage

Learn more about physical storage, logical volumes, and the devices and configurations that are supported by the Virtual I/O Server.

Physical volumes

Physical volumes can be exported to client partitions as virtual Small Computer Serial Interface (SCSI) disks. The Virtual I/O Server (VIOS) is capable of taking a pool of heterogeneous physical disk storage attached to its backend and exporting this as homogeneous storage in the form of SCSI disk LUNs.

The VIOS must be able to accurately identify a physical volume each time it boots, even if an event such as a storage area network (SAN) reconfiguration or adapter change has taken place. Physical volume attributes, such as the name, address, and location, might change after the system reboots due to SAN reconfiguration. However, the VIOS must be able to recognize that this is the same device and update the virtual device mappings. Hence, to export a physical volume as a virtual device, the physical volume must have either a unique identifier (UDID), a physical identifier (PVID), or an IEEE volume attribute.

For instructions about determine whether your disks have one of these identifiers, see "Identifying exportable disks" on page 120.

The following commands are used to manage physical volumes.

Table 9. Physical volume commands and their descriptions

Physical volume command	Description
lspv	Displays information about physical volumes within the VIOS logical partition.
migratepv	Moves allocated physical partitions from one physical volume to one or more other physical volumes.

Logical volumes

Understand how logical volumes can be exported to client partitions as virtual Small Computer Serial Interface (SCSI) disks. A logical volume is a portion of a physical volume.

A hierarchy of structures is used to manage disk storage. Each individual disk drive or LUN, called a *physical volume*, has a name, such as **/dev/hdisk0**. Every physical volume in use either belongs to a volume group or is used directly for virtual storage. All of the physical volumes in a volume group are divided into physical partitions of the same size. The number of physical partitions in each region varies, depending on the total capacity of the disk drive.

Within each volume group, one or more logical volumes are defined. Logical volumes are groups of information that is located on physical volumes. Data on logical volumes appears to the user to be contiguous but can be discontiguous on the physical volume. This allows logical volumes to be resized or relocated and to have their contents replicated.

Each logical volume consists of one or more logical partitions. Each logical partition corresponds to at least one physical partition. Although the logical partitions are numbered consecutively, the underlying physical partitions are not necessarily consecutive or contiguous.

After installation, the system has one volume group (the rootvg volume group) consisting of a base set of logical volumes that are required to start the system.

You can use the commands described in the following table to manage logical volumes.

Table 10. Logical volume commands and their descriptions

Logical volume command	Description
chlv	Changes the characteristics of a logical volume.
cplv	Copies the contents of a logical volume to a new logical volume.
extendlv	Increases the size of a logical volume.
lslv	Displays information about the logical volume.
mklv	Creates a logical volume.
mklvcopy	Creates a copy of a logical volume.
rmlv	Removes logical volumes from a volume group.
rmlvcopy	Removes a copy of a logical volume.

Creating one or more distinct volume groups rather than using logical volumes that are created in the rootvg volume group allows you to install any newer versions of the Virtual I/O Server while maintaining client data by exporting and importing the volume groups created for virtual I/O.

Notes:

- Logical volumes used as virtual disks must be less than one TB (where TB equals 1 099 511 627 776 bytes) in size.
- For best performance, avoid using logical volumes (on the Virtual I/O Server) as virtual disks that are mirrored or striped across multiple physical volumes.

Volume groups

Find information about volume groups.

A volume group is a type of storage pool that contains one or more physical volumes of varying sizes and types. A physical volume can belong to only one volume group per system. There can be up to 4096 active volume groups on the Virtual I/O Server.

When a physical volume is assigned to a volume group, the physical blocks of storage media on it are organized into physical partitions of a size determined by the system when you create the volume group. For more information, see "Physical partitions" on page 19.

When you install the Virtual I/O Server, the root volume group called rootvg is automatically created that contains the base set of logical volumes required to start the system logical partition. The rootvg includes paging space, the journal log, boot data, and dump storage, each in its own separate logical volume. The rootvg has attributes that differ from user-defined volume groups. For example, the rootvg cannot be imported or exported. When you use a command or procedure on the rootvg, you must be familiar with its unique characteristics.

Table 11. Frequently used volume group commands and their descriptions

Command	Description
activatevg	Activates a volume group
chvg	Changes the attributes of a volume group
deactivatevg	Deactivates a volume group
exportvg	Exports the definition of a volume group
extendvg	Adds a physical volume to a volume group
importvg	Imports a new volume group definition

Table 11. Frequently used volume group commands and their descriptions (continued)

Command	Description
lsvg	Displays information about a volume group
mkvg	Creates a volume group
reducevg	Removes a physical volume from a volume group
syncvg	Synchronizes logical volume copies that are not current

Small systems might require only one volume group to contain all of the physical volumes (beyond the rootvg volume group). You can create separate volume groups to make maintenance easier because groups other than the one being serviced can remain active. Because the rootvg must always be online, it contains only the minimum number of physical volumes necessary for system operation. It is suggested that the rootvg not be used for client data.

You can move data from one physical volume to other physical volumes in the same volume group by using the **migratepv** command. This command allows you to free a physical volume so it can be removed from the volume group. For example, you could move data from a physical volume that is to be replaced.

Physical partitions

This topic contains information about physical partitions.

When you add a physical volume to a volume group, the physical volume is partitioned into contiguous, equal-sized units of space called *physical partitions*. A physical partition is the smallest unit of storage space allocation and is a contiguous space on a physical volume.

Physical volumes inherit the volume group's physical partition size.

Logical partitions

This topic contains information logical storage partitions.

When you create a logical volume, you specify its size in megabytes or gigabytes. The system allocates the number of logical partitions that are required to create a logical volume of at least the specified size. A logical partition is 1 or 2 physical partitions, depending on whether the logical volume is defined with mirroring enabled. If mirroring is disabled, there is only one copy of the logical volume (the default). In this case, there is a direct mapping of one logical partition to one physical partition. Each instance, including the first, is called a copy.

Quorums

Find information about quorums.

A quorum exists when most of Volume Group Descriptor Areas and Volume Group Status Areas (VGDA/VGSA) and their disks are active. A quorum ensures data integrity of the VGDA/VGSA in the event of a disk failure. Each physical disk in a volume group has at least one VGDA/VGSA. When a volume group is created onto a single disk, the volume group initially has two VGDA/VGSA on the disk. If a volume group consists of two disks, one disk still has two VGDA/VGSA, but the other disk has one VGDA/VGSA. When the volume group is made up of three or more disks, each disk is allocated just one VGDA/VGSA.

A quorum is lost when enough disks and their VGDA/VGSA are unreachable so that a 51% majority of VGDA/VGSA no longer exists.

When a quorum is lost, the volume group deactivates itself so that the disks are no longer accessible by the logical volume manager. This prevents further disk I/O to that volume group so that data is not lost or assumed to be written when physical problems occur. As a result of the deactivation, the user is notified in the error log that a hardware error has occurred and service must be performed.

A volume group that has been deactivated because its quorum has been lost can be reactivated by using the **activatevg -f** command.

Virtual media repository

The virtual media repository provides a single container to store and manage file-backed virtual optical media files. Media stored in the repository can be loaded into file-backed virtual optical devices for exporting to client partitions.

Only one repository can be created within a Virtual I/O Server.

The virtual media repository is available with Virtual I/O Server Version 1.5, or later.

The virtual media repository is created and managed by using the following commands.

Table 12. Virtual media repository commands and their descriptions

Command	Description
chrep	Changes the characteristics of the virtual media repository
chvopt	Changes the characteristics of a virtual optical media
loadopt	Loads file-backed virtual optical media into a file-backed virtual optical device
lsrep	Displays information about the virtual media repository
lsvopt	Displays information about file-backed virtual optical devices
mkrep	Creates the virtual media repository
mkvdev	Creates file-backed virtual optical devices
mkvopt	Creates file-backed virtual optical media
rmrep	Removes the virtual media repository
rmvopt	Removes file-backed virtual optical media
unloadopt	Unloads file-backed virtual optical media from a file-backed virtual optical device

Optical devices

Optical devices can be exported by the Virtual I/O Server. This topic gives information about what types of optical devices are supported.

The Virtual I/O Server supports exporting optical Small Computer Serial Interface (SCSI) devices. These are referred to as *virtual SCSI optical devices*. Virtual optical devices can be backed by DVD drives or files. Depending on the backing device, the Virtual I/O Server exports a virtual optical device with one of following profiles:

- DVD-ROM
- DVD-RAM

Virtual optical devices that are backed by physical optical devices can be assigned to only one client logical partition at a time. To use the device on a different client logical partition, it must first be removed from its current logical partition and reassigned to the logical partition that uses the device.

Tape

Tape devices can be exported by the Virtual I/O Server. This topic gives information about what types of tape devices are supported.

The Virtual I/O Server supports exporting physical tape devices to client logical partitions. These are referred to as *virtual Small Computer Serial Interface (SCSI) tape devices*. Virtual SCSI tape devices are backed up by physical tape devices.

Virtual SCSI tape devices are assigned to only one client logical partition at any given time. To use the device on a different client logical partition, it must first be removed from its current logical partition and reassigned to the logical partition that uses the device.

Restriction:

- The physical tape device must be attached by a serial-attached SCSI (SAS) or Universal Serial Bus (USB) tape device and both the drive types must be DAT320.
- The Virtual I/O Server does not support media movers, even if the physical device supports them.
- It is suggested that you assign the tape device to its own Virtual I/O Server adapter because as tape devices often send large amounts of data, which might affect the performance of any other device on the adapter.

Virtual storage

Disks, tapes, Universal Serial Bus (USB) mass storage, and optical devices are supported as virtual Small Computer Serial Interface (SCSI) devices. This topic describes how those devices function in a virtualized environment and provides information on what devices are supported.

The Virtual I/O Server might virtualize or export, disks, tapes, USB mass storage, and optical devices, such as CD-ROM drives and DVD drives, as virtual devices. For a list of supported disks and optical devices, see the data sheet available on the Fix Central website. For information about configuring virtual SCSI devices, see "Creating the virtual target device on the Virtual I/O Server " on page 109.

Disk

Disk devices can be exported by the Virtual I/O Server. This topic gives information about what types of disks and configurations are supported.

The Virtual I/O Server supports exporting disk Small Computer Serial Interface (SCSI) devices. These are referred to as *virtual SCSI disks*. All virtual SCSI disks must be backed by physical storage. The following types of physical storage can be used to back virtual disks:

- Virtual SCSI disk backed by a physical disk
- Virtual SCSI disk backed by a logical volume
- Virtual SCSI disk backed by a file

Regardless of whether the virtual SCSI disk is backed by a physical disk, logical volume, or a file, all standard SCSI rules apply to the device. The virtual SCSI device behaves as a standard SCSI-compliant disk device, and it can serve as a boot device or a Network Installation Management (NIM) target, for example.

Virtual SCSI Client Adapter Path Timeout

The virtual SCSI Client Adapter Path Timeout feature allows the client adapter to detect whether a Virtual I/O Server is not responding to I/O requests. Use this feature only in configurations in which devices are available to a client logical partition from multiple **Virtual I/O Servers**. These configurations could be one of the following:

- Multipath I/O (MPIO) configurations
- Configurations where a volume group is mirrored by devices on multiple **Virtual I/O Servers**.

vSCSI client adapter path timeout scenarios

If no I/O requests issued to the virtual SCSI server adapter are serviced within the number of seconds specified by the virtual SCSI path timeout value, one more attempt is made to contact the virtual SCSI server adapter, waiting up to 60 seconds for a response.

If, after 60 seconds, there is still no response from the server adapter, all outstanding I/O requests to that adapter fail and an error is written to the client logical partition error log.

- If MPIO is being used, the MPIO Path Control Module retries the I/O requests on another path. Otherwise, the failed requests are returned to the applications.
- If the devices on this adapter are part of a mirrored volume group, those devices are marked as *missing* and the Logical Volume Manager logs errors in the client logical partition error log.

If one of the failed devices is the root volume group (rootvg) for the logical partition, and the rootvg is not available through another path or is not being mirrored on another Virtual I/O Server, the client logical partition is likely to shut down. The virtual SCSI client adapter attempts to reestablish communication with the Virtual I/O Server and logs a message in the system error log when it is able to do so. Mirrored volume groups must be manually resynchronized by running the **varyonvg** command when the missing devices are once again available.

A configurable virtual SCSI client adapter ODM attribute, **vscsi_path_to**, is provided. This is a tunable attribute that is specific to an AIX client. The path timeouts for the Linux operating system are configured differently. This attribute is used both to indicate whether the feature is enabled and to store the value of the path timeout, if the feature is enabled.

The system administrator sets the ODM attribute to 0 to disable the feature, or to the time, in seconds, to wait before checking if the path to the server adapter has failed. If the feature is enabled, a minimum setting of 30 seconds is required. If a setting is entered between 0 and 30 seconds, the value is changed to 30 seconds upon the next adapter reconfiguration or reboot.

This feature is disabled by default, thus the default value of **vscsi_path_to** is 0. You must exercise careful consideration when setting this value, keeping in mind that when the virtual SCSI server adapter is servicing the I/O request, the storage device the request is being sent to might be either local to the Virtual I/O Server or on a SAN.

The **vscsi_path_to** client adapter attribute can be set by using the SMIT utility or by using the **chdev -P** command. The attribute setting can also be viewed by using SMIT or the **lsattr** command. The setting does not take effect until the adapter is reconfigured or the client partition is rebooted.

Virtual SCSI device read or write command timeout

The virtual SCSI device read or write command timeout feature facilitates the virtual SCSI device to detect a hung I/O request. You can use this feature in any virtual SCSI client configuration to detect and recover from the I/O request failures. The following configurations are supported:

- Virtual SCSI clients in which disks are exported through a single virtual SCSI server adapter.
- Same disks are available to the virtual SCSI clients from multiple virtual SCSI server adapters.

If the virtual SCSI device read or write command timeout feature is enabled, all the read or write command requests that are issued to the virtual SCSI server adapter are timed. If any read or write command is not serviced within the number of seconds that is specified by the command timeout value, then the virtual SCSI client adapter causes the command to time-out. The connection with the virtual SCSI server adapter is then closed and subsequently, a new connection is reinitialized.

A configurable virtual SCSI device ODM attribute, **rw_timeout** is specified. This attribute is a tunable attribute and indicates the read or write command timeout value for the device that is configured on the virtual SCSI client. You can modify the **rw_timeout** attribute for the virtual SCSI device by using the **chdev** or **chdev -P** command. You can use the **lsattr -R -l device -a rw_timeout** command that provides the range of values that can be used for the device read or write command timeout feature. You must specify the value for the read or write command timeout feature within the range of values indicated by the **lsattr -R -l device -a rw_timeout** command. If the specified value for the read or write command timeout feature is less than the minimum or greater than the maximum value indicated in the range of values, the **chdev** command returns an error.

The read or write command timeout feature is enabled by default from AIX 7.2 TL 2, AIX 7.1 TL 5, and later. This feature is disabled in the earlier AIX releases, by default.

The **rw_timeout** attribute is associated with every virtual SCSI device and not just the disk. With AIX 7.2 TL 5, and later, the **rw_timeout** attribute is not a virtual SCSI client adapter attribute.

The following table provides details about the default range and acceptable range (in seconds) of the read or write command timeout value.

Table 13. Default and acceptable range (in seconds) of the read or write command timeout value

AIX release	Default state	Default value	Minimum value	Maximum value
AIX 7.2 TL 5, and later	Enabled	45	device-specific	device-specific
AIX 7.2 TL 2, AIX 7.1 TL 5, and later	Enabled	45	45	3600
AIX 7.2 TL 1, AIX 7.1 TL 4, and other	Disabled	0	120	3600

Optical

Optical devices can be exported by the Virtual I/O Server. This topic gives information about what types of optical devices are supported.

The Virtual I/O Server supports exporting physical optical devices to client logical partitions. These are referred to as *virtual Small Computer Serial Interface (SCSI) optical devices*. Virtual SCSI optical devices can be backed by DVD drives or files. Depending on the backing device, the Virtual I/O Server exports a virtual optical device with one of following profiles:

- DVD-ROM
- DVD-RAM

For example, file-backed virtual SCSI optical devices are exported as DVD-RAM devices. File-backed virtual SCSI optical devices can be backed by read/write or read-only files. Depending on the file permissions, the device can appear to contain a DVD-ROM or DVD-RAM disk. Read/write media files (DVD-RAM) cannot be loaded into more than one file-backed virtual SCSI optical device simultaneously. Read-only media files (DVD-ROM) can be loaded into multiple file-backed virtual SCSI optical devices simultaneously.

Virtual SCSI optical devices that are backed by physical optical devices can be assigned to only one client logical partition at any given time. To use the device on a different client logical partition, it must first be removed from its current logical partition and reassigned to the logical partition that uses the device.

Virtual SCSI optical devices always appear as SCSI devices on the client logical partitions regardless of whether the device type exported from the Virtual I/O Server is a SCSI, IDE, USB device, or a file.

Tape

Tape devices can be exported by the Virtual I/O Server. This topic gives information about what types of tape devices are supported.

The Virtual I/O Server supports exporting physical tape devices to client logical partitions. These are referred to as *virtual Small Computer Serial Interface (SCSI) tape devices*. Virtual SCSI tape devices are backed up by physical tape devices.

Virtual SCSI tape devices are assigned to only one client logical partition at any given time. To use the device on a different client logical partition, it must first be removed from its current logical partition and reassigned to the logical partition that uses the device.

Restriction:

- The physical tape device must be attached by a serial-attached SCSI (SAS) or Universal Serial Bus (USB) tape device and both the drive types must be DAT320.
- The Virtual I/O Server does not support media movers, even if the physical device supports them.
- It is suggested that you assign the tape device to its own Virtual I/O Server adapter because as tape devices often send large amounts of data, which might affect the performance of any other device on the adapter.

USB mass storage

Universal Serial Bus (USB) mass storage devices are exported by the Virtual I/O Server. This topic gives information about the types of supported USB devices and configurations.

The Virtual I/O Server exports the USB attached hard disk devices to the client logical partitions. These exported devices are referred to as *virtual Small Computer System Interface (SCSI) USB disk devices*. The virtual SCSI USB disk devices are backed up by the physical USB mass storage devices. The virtual SCSI USB disk is used to back up or restore data of the client logical partitions. These disks can also be used as a boot device.

The virtual SCSI USB disk devices are assigned to only one client logical partition at any given time. To use the device on a different client logical partition, it must first be removed from its current logical partition and then reassigned to the logical partition that uses the device.

Device compatibility in a Virtual I/O Server environment

Learn more about virtual-to-physical device compatibility in a Virtual I/O Server environment.

The virtual-to-physical device (p2v) compatibility that is described in this topic refers only to the data on the device, not necessarily to the capabilities of the device. A device is p2v compatible when the data retrieved from that device is identical regardless of whether it is accessed directly through a physical attachment or virtually (for example, through the Virtual I/O Server). That is, every logical block (for example, LBA 0 through LBA n-1) returns identical data for both physical and virtual devices. Device capacity must also be equal to claim p2v compliance. You can use the Virtual I/O Server **chkdev** command to determine if a device is p2v compatible.

Virtual disk devices exported by the Virtual I/O Server are referred to as virtual Small Computer Serial Interface (SCSI) disks. A virtual SCSI disk device might be backed by an entire physical volume, a logical volume, a multi-path device, or a file.

Data replication (such as copy services) and device movement between physical and virtual environments are common operations in today's data center. These operations, involving devices in a virtualized environment, often have a dependency on p2v compliance.

Copy Services refer to various solutions that provide data replication function including data migration, flashcopy, point-in-time copy, and remote mirror and copy solutions. These capabilities are commonly used for disaster recovery, cloning, backup/restore, and more.

Device movement between physical and virtual environments refers to the ability to move a disk device between physical (for example, a directly attached SAN) and virtual I/O (for example, Virtual I/O Server that is attached to a SAN) environments and use the disk without having to back up or restore the data. This capability is useful for server consolidation.

The operations might work if the device is p2v compatible. However, not all device combinations and data replication solutions have been tested by IBM. See claims by the Copy Services vendor for support claims for devices managed by Virtual I/O Server.

A device is p2v compatible if it meets the following criteria:

- It is an entire physical volume (for example, a LUN)
- Device capacity is identical in both physical and virtual environments
- The Virtual I/O Server is able to manage this physical volume by using a UDID or iEEE ID.

Devices managed by the following multipathing solutions within the Virtual I/O Server are expected to be UDID devices.

- All multipath I/O (MPIO) versions, including Subsystem Device Driver Path Control Module (SDDPCM), EMC PCM, and Hitachi Dynamic Link Manager (HDLM) PCM
- EMC PowerPath 4.4.2.2 or later
- IBM Subsystem Device Driver (SDD) 1.6.2.3 or later
- Hitachi HDLM 5.6.1 or later

Virtual SCSI devices created with earlier versions of PowerPath, HDLM, and SDD are not managed by UDID format and are not expected to be p2v compliant. The operations mentioned, such as data replication or movement between Virtual I/O Server and non-Virtual I/O Server environments) are not likely to work in these cases.

Related tasks

Determining whether a physical volume is managed by UDID or IEEE
Determine whether a physical volume is or can be managed by a unit device identifier (UDID) or IEEE. You can use the Virtual I/O Server **chkdev** command to display this data.

Related information

chkdev command

Determining whether a physical volume is managed by UDID or IEEE
Determine whether a physical volume is or can be managed by a unit device identifier (UDID) or IEEE. You can use the Virtual I/O Server **chkdev** command to display this data.

Before you begin

To determine whether a physical volume is or can be managed by the UDID format, the following must be verified:

- If it is an existing Virtual I/O Server LUN, determine whether its format is UDID.
- If it is a LUN to be moved to Virtual I/O Server, first verify that the Virtual I/O Server is prepared to see that LUN as a UDID LUN, by checking it at the source host.

  Note: Moving a physical disk to a Virtual I/O Server that is not capable of managing the device by using UDID might result in data loss. In this case, back up the data before allocating the LUN to the Virtual I/O Server.

Procedure

1. To determine whether a device has a UDID or an IEEE volume attribute identifier for the Virtual I/O Server, type:

   ```
   chkdev -verbose
   ```

   Output similar to the following example is displayed:

   ```
   NAME:                hdisk1
   IDENTIFIER:          210ChpO-c4HkKBc904N37006NETAPPfcp
   PHYS2VIRT_CAPABLE:   YES
   VIRT2NPIV_CAPABLE:   NA
   VIRT2PHYS_CAPABLE:   NA
   PVID:                00c58e40599f2f900000000000000000
   UDID:                2708ECVBZ1SC10IC35L146UCDY10-003IBXscsi
   IEEE:
   VTD:

   NAME:                hdisk2
   IDENTIFIER:          600A0B800012DD0D00000AB441ED6AC
   PHYS2VIRT_CAPABLE:   YES
   VIRT2NPIV_CAPABLE:   NA
   VIRT2PHYS_CAPABLE:   NA
   PVID:                00c58e40dcf83c850000000000000000
   UDID:
   IEEE:                600A0B800012DD0D00000AB441ED6AC
   VTD:
   ```

   If the *IEEE:* field does not appear, then the device does not have an IEEE volume attribute identifier.
2. To determine whether a device has an UDID for the AIX operating system, type:

   ```
   odmget -qattribute=unique_id CuAt
   ```

   The disks that have a UDID are listed. Output similar to the following example is displayed:

```
CuAt:
    name = "hdisk1"
    attribute = "unique_id"
    value = "2708ECVBZ1SC10IC35L146UCDY10-003IBXscsi"
    type = "R"
    generic = ""
    rep = "nl"
    nls_index = 79

CuAt:
    name = "hdisk2"
    attribute = "unique_id"
    value = "210800038FB50AST373453LC03IBXscsi"
    type = "R"
    generic = ""
    rep = "nl"
nls_index = 79
```

3. To determine whether a device has an UDID for the AIX operating system, type:

```
odmget -qattribute=unique_id CuAt
```

The disks that have a UDID are listed. Output similar to the following example is displayed:

```
CuAt:
    name = "hdisk1"
    attribute = "unique_id"
    value = "2708ECVBZ1SC10IC35L146UCDY10-003IBXscsi"
    type = "R"
    generic = ""
    rep = "nl"
    nls_index = 79

CuAt:
    name = "hdisk2"
    attribute = "unique_id"
    value = "210800038FB50AST373453LC03IBXscsi"
    type = "R"
    generic = ""
    rep = "nl"
nls_index = 79
```

4. To determine whether a device has an IEEE volume attribute identifier for the AIX operating system, type:

```
lsattr -l hdiskX
```

Disks with an IEEE volume attribute identifier have a value in the *ieee_volname* field. Output similar to the following example is displayed:

```
...
cache_method   fast_write                       Write Caching method
ieee_volname   600A0B800012DD0D00000AB441ED6AC IEEE Unique volume name
lun_id         0x001a000000000000               Logical Unit Number
        ...
```

If the *ieee_volname* field does not appear, then the device does not have an IEEE volume attribute identifier.

Note: DS4K and FAStT storage that use the Redundant Disk Array Controller (RDAC) driver for multipathing are managed by using an IEEE ID.

5. To determine whether a device has an IEEE volume attribute identifier for the AIX operating system, type:

```
lsattr -l hdiskX
```

Disks with an IEEE volume attribute identifier have a value in the *ieee_volname* field. Output similar to the following example is displayed:

```
...
cache_method   fast_write                       Write Caching method
ieee_volname   600A0B800012DD0D00000AB441ED6AC IEEE Unique volume name
```

```
lun_id         0x001a000000000000              Logical Unit Number
        ...
```

If the *ieee_volname* field does not appear, then the device does not have an IEEE volume attribute identifier.

Note: DS4K and FAStT storage that use the Redundant Disk Array Controller (RDAC) driver for multipathing are managed by using an IEEE ID.

Related information

chkdev command

Cache device management

Learn about cache device management in a Virtual I/O Server (VIOS) environment.

The cache device management feature creates an infrastructure to manage attached solid-state drives (SSDs) for caching on client partitions.

Note: Cache engine is not available on the VIOS. Caching of target devices on VIOS is not supported.

Cache management concepts

Cache device
Is the SSD or flash disk device used for caching.

Cache pool
Is a group of cache devices that is only used for disk caching. A cache pool (or volume group) provides a simplified way to manage multiple flash disk devices. You can add additional devices to expand a cache pool, as needed. Currently, only a single cache pool is supported.

Cache partition
Is a logical cache device that is created out of a cache pool. A cache partition (or logical volume) provides flexibility and better utilization of flash storage for caching. It allows you to use multiple partitions / logical cache devices. Partitions can be expanded as needed for a larger working set. A cache partition must be assigned to a virtual SCSI server adapter.

The cache_mgt command provides the infrastructure that is required to manage caching on solid state drive (SSD) devices.

For more information, see the cache_mgt Command.

Mapping devices

Mapping devices are used to facilitate the mapping of physical resources to a virtual device.

iSCSI disk support for VIOS

The Internet Small Computer Systems Interface (iSCSI) disk is supported in the Virtual I/O Server (VIOS) 3.1.0, or later, and requires FW 860.20, or later. The FW level of 860.20 is supported on POWER8 processor-based systems. For a POWER9 processor-based systems, the minimum FW level required is FW 910.

The Internet Small Computer Systems Interface (iSCSI) disk provides block-level access to storage devices by carrying SCSI commands over an Internet Protocol network. The *iSCSI* disk is used to facilitate data transfers over the internet by using TCP, a reliable transport mechanism that uses either IPV6 or IPV4 protocols. The *iSCSI* disk is used to manage storage over long distances.

The *iSCSI* support in VIOS allows *iSCSI* disks to be exported to client logical partitions as virtual disks (vSCSI disks). This support is available in VIOS version 3.1, and later, on both POWER8 and POWER9 systems. If you are using a POWER8 system, the firmware level must be at FW860.20 or later. There are no minimum firmware level requirements for POWER9 systems. POWER9 systems can run on various firmware levels like FW910, FW920, FW930, or later.

VIOS version 3.1 enables Multipath I/O (MPIO) support for the *iSCSI* initiator. With MPIO support, you can configure and create multiple paths to an *iSCSI* disk, similar to other protocols. The client logical partition can run either an AIX or Linux operating system.

VIOS version 3.1.1 enables support for multiple *iSCSI* initiators on the VIOS. This support also includes performance enhancements for the *iSCSI* driver. With multiple *iSCSI* initiator support, you can create multiple *iSCSI* software initiator devices on a single AIX operating system instance.

The advantages of configuring multiple *iSCSI* software initiators are as follows:

- You can easily create multiple paths for an *iSCSI* disk that supports Multipath I/O (MPIO). Each path creates its own TCP/IP socket connection. Thereby, the iSCSI traffic is spread across more connections to improve performance through increased concurrent processing.
- Multiple I/O requests from the *iSCSI* disk can be logically separated. This reduces the chances of I/O request conflicts between applications.

Limitations

Currently, the *iSCSI* disk support for VIOS has the following limitations:

- There is no VIOS boot support using an *iSCSI* disk.
- The flat file-based discovery policy is not supported.
- The *iSCSI* disk based logical volume (LV) backed devices are not supported.
- Shared Storage Pools using *iSCSI* disks as either Repo or Shared Pool disks is not supported.
- The *iSCSI* disks or *iSCSI* based LVs or volume groups (VGs) cannot be used as paging devices for the Active Memory Sharing (AMS) or Remote restart feature.
- If the backing device is an *iSCSI* disk, the `client_reserve` and `mirrored` attribute are not supported for virtual target devices.
- On VIOS version 3.1, booting from an *iSCSI* disk is not supported.

Recommendations

For optimal performance of the *iSCSI* disk, the following hardware configuration is recommended.

- A separate private network to access the *iSCSI* storage.
- Use of high-speed network adapters and switches (at least 10G is recommended).

Related reference

iSCSI software initiator and software target

Related information

chiscsi command

lsiscsi command

mkiscsi command

rmiscsi command

Shared storage pools

Learn about shared storage pools on the Virtual I/O Server.

Clusters

Learn about using the Virtual I/O Server (VIOS) and creating a clustering configuration.

The following table provides details about the number of VIOS partitions allowed in a cluster, in different VIOS versions.

Table 14. Allowed VIOS partitions in a cluster

VIOS version	Allowed VIOS partitions in a cluster
VIOS 2.2.0.11, Fix Pack 24, Service Pack 1	1
VIOS 2.2.2.0, or later	16

Thus, a cluster consists of a up to 16 VIOS logical partitions with a shared storage pool that provides distributed storage access to the VIOS logical partitions in the cluster. Each cluster requires a separate repository disk and shared storage pool disks. The shared storage pool can be accessed by all VIOS logical partitions in the cluster.

All the The VIOS logical partitions within a cluster must have access to all the physical volumes in a shared storage pool.

You can create and manage clusters by using the commands in the following table.

Table 15. Cluster commands and their descriptions

Command	Description
cluster	Provides cluster management and listing capabilities.
chrepos	Replaces the repository disk.

The following table lists the scalability limits for clusters in the VIOS Version 2.2.2.0, or later:

Table 16. Scalability limits for clusters

Component	Minimum value	Maximum value
Number of VIOS systems in a cluster	1	16
Number of physical disks in the shared storage pool	1	1024
Number of logical unit mappings in the shared storage pool	1	8192
Number of client logical partitions per VIOS	1	250
Storage capacity of physical disks in the shared storage pool	5 GB	16 TB
Storage capacity of the shared storage	5 GB	512 TB
Storage capacity of a logical unit in the shared storage	1 GB	4 TB
Number of repository disks	1	1
Mirror copies	1	2
Number of mirror copies per shared storage pool	1	2

Related tasks

Replacing a repository disk

On the Virtual I/O Server (VIOS) Version 2.2.2.0, you can replace a repository disk by using the VIOS command-line interface.

Storage pools

Learn about logical volume storage pools and file storage pools.

The following table lists the various types of storage pools.

Table 17. Storage pools

Storage pools supported	Virtual I/O Server (VIOS) release
• Logical volume storage pools (LVPOOL) • File storage pools (FBPOOL)	VIOS Version 1.5, and later
Shared storage pools	VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, and later

Like volume groups, logical volume storage pools are collections of one or more physical volumes. The physical volumes that comprise a logical volume storage pool can be of varying sizes and types. File storage pools are created within a parent logical volume storage pool and contain a logical volume that contains a file system with files.

Logical volume storage pools store logical volume backing devices, file-backed storage pools, and the virtual media repository. File storage pools store file-backing devices.

For using storage pools, it is not necessary for you to have extensive knowledge of how to manage volume groups and logical volumes to create and assign logical storage to a client logical partition. Devices that are created by using a storage pool are not limited to the size of the individual physical volumes.

On the VIOS, you can use shared storage pools. Shared storage pools provide distributed storage access to all the VIOS logical partitions in a cluster.

Storage pools are created and managed by using the following commands.

Table 18. Storage pool commands and their descriptions

Command	Description
alert	Sets, removes, and lists all the alerts for the storage pool in a cluster.
chsp	Changes the characteristics of a storage pool.
chbdsp	Changes the characteristics of a backing device within a storage pool.
failgrp	Manages mirroring in storage pools.
lu	Manages logical units in shared storage pools.
lssp	Displays information about a storage pool.
mkbdsp	Assigns storage from a storage pool to be a backing device for a virtual Small Computer Serial Interface (SCSI) adapter.
mksp	Creates a storage pool. This storage pool is created by default when you create a cluster.
pv	Manages physical storage in shared storage pools.
rmbdsp	Removes a backing device from its virtual SCSI adapter, or a VIOS object (Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later), and returns the storage back to the storage pool.
rmsp	Removes a file storage pool. This storage pool is removed by default when you remove a cluster.

Table 18. Storage pool commands and their descriptions (continued)

Command	Description
snapshot	Creates, deletes, and rolls back a snapshot image of a single logical unit or multiple logical units.
tier	Manages storage tiers in a shared storage pool.

In VIOS logical partitions prior to Version 2.2.0.11, Fix Pack 24, Service Pack 1, each VIOS logical partition has a single default storage pool that can be modified only by the prime administrator. By default, *rootvg*, which is a logical volume pool, is the default storage pool unless the prime administrator configures a different default storage pool.

Do not create client storage in rootvg. By creating one or more distinct logical volume storage pools rather than using the rootvg volume group, you can install any newer versions of the VIOS while maintaining client data by exporting and importing the volume groups created for virtual I/O.

Unless explicitly specified otherwise, the storage pool commands operate on the default storage pool. This situation can be useful on systems that contain most or all of its backing devices in a single storage pool.

Note: A physical volume can be assigned only to one virtual function at a time. For example, a physical volume that is used by a storage pool cannot be assigned for use as a virtual disk at the same time.

Storage tiers

Storage tiers allow you to group physical volumes (PVs) within a storage pool.

Storage tiers provide flexibility to group disks in ways that can improve the management of your environment. Some possible improvements that storage tiers can provide are identified in the following list:

- Data security: You can group disks into security classes. For example, you can group one set of disks in an ultra-secure room and another set of disks with simple encryption in a less-secure location.
- Performance: You can group disks by I/O speed. By grouping this way, you can ensure that your most frequently accessed information is on your fastest storage media.
- Reliability: You can isolate storage pool metadata from user data. This helps to increase the reliability of the system because it is not simultaneously accessing the different types of information on the same disk. Storage tiers also help reliability by enabling extra mirroring of meta/critical data.

The initial storage tier that is created when you create a cluster is called the *system tier*. It is automatically given the name of SYSTEM. All operations happen in this storage tier by default until you create a different storage tier and identify it as the default tier. Pool metadata and file metadata are always maintained in the system tier. This storage tier is sometimes referred to as *tier 0*.

There are two types of system tiers that you can configure. The type is determined by the data that is contained within the system tier. The default configuration is the co-mingled (or *unrestricted*) storage tier, which contains both metadata and user data. You can change the system tier to a *restricted* storage tier, which contains only metadata. Restricted tiers do not allow user data, so you must create a *user tier* to store your user data (logical units).

You can create additional *user tiers*. The total number of tiers allowed is 10, including the system tier.

If you created a cluster with an unrestricted tier but decide to isolate the pool metadata, you can set the unrestricted system tier as a restricted system tier. When making a system tier into a restricted tier, you do not have to move all of the existing user data from the system tier. Any LUs that are in the system tier when you restrict it remain there until you are ready to move them. Note that for thin-provisioned LUs still assigned to the restricted system tier, the storage pool places new block allocations in the system tier.

Logical units (LUs) can be assigned to an unrestricted system tier or to a specific user tier. During the creation of LUs, they get assigned to the tiers that are specified in the **-tier** option.

With the addition of multiple-tier support, administrators are allowed to do the following:

- Create LUs in specific tiers
- Add a PV to a specific tier
- Remove a PV from a tier
- Create failure groups within tiers
- Move LUs between tiers
- Create new tiers
- Remove tiers (except the pool's *system* tier, which can be removed only when the cluster is deleted)
- List tiers and also provide details about a specific tier.
- List PVs in a tier
- List LUs in a tier
- Monitor individual tiers

Related concepts

Managing storage tiers
You can use the command-line interface on the Virtual I/O Server (VIOS) to manage a storage tier. You can also use the Hardware Management Console (HMC) version 8.4.0, or later to manage storage tiers.

Related information

tier command

Failure group

Shared Storage Pool (SSP) Mirroring is enabled from Virtual I/O Server (VIOS) Version 2.2.3. Mirroring an SSP is an optional step that increases resiliency by adding redundancy. Inside the storage pool, there might be two sets of shared logical unit numbers (LUNs, or physical volumes (PVs)). These two named sets of LUNs are referred to as *failure groups* or *mirrors*. The preferred practice is to define the two failure groups on different physical storage arrays for best availability.

The whole pool is either a single copy pool (one failure group) or double copy pool (two failure groups). If two groups are defined, the whole pool is mirrored and not just individual logical units (LUs) of PVs. The data space that belongs to an LU is divided into 64 MB each and the LUs are placed in individual physical volumes (LUNs) in the pool. The exact data placement is decided in the background. Hence, it is not an exact one-to-one mirroring.

By default, a single copy pool is created by running the **`cluster -create`** command and the first failure group is named `Default`. You can rename the first failure group and add a second failure group.

Consider the following characteristics of a mirrored SSP:

- A mirrored SSP doubles the disk space requirement, which is typical for Disaster Recovery (DR) solutions.
- A mirrored SSP is completely transparent for client VMs. Therefore, there is no action needed on the client operating system. The VIOS accesses the storage and keeps the mirrors in a synchronized state. The VIOS creates duplicates, writes to both mirrors and performs re-mirroring if one of the mirrors becomes out-of-sync.
- The VIOS performs recovery and re-mirroring in the background, without affecting the client VMs.

The following preferred practices relate to mirrored storage pools:

- Failure groups must be of the same size. If there are two failure groups in an SSP and their capacity is not the same, the total size of the SSP available for allocation of LUs is the sum of the capacity of LUNs that are in the smaller failure group. The rest of the capacity in the larger failure group is not used.
- When you create a large mirrored pool with two failure groups, the preferred practice is to create a pool of one disk and add the second failure group to mirror the first pool. Then, you can add physical volumes to both failure groups to increase the capacity of the pool.
- If a disk or a storage controller in a single failure group fails, the mirrored storage pool runs in a degraded state. In this case, you must take corrective actions to resolve the issue on the storage controller.

- The system firmware must be upgraded to the latest release to achieve the optimum performance from the mirrored storage pools.

Related information

failgrp command

Virtual networking

Learn about virtual Ethernet, Host Ethernet Adapter (or Integrated Virtual Ethernet), Internet Protocol version 6 (IPv6), Link Aggregation (or Etherchannel), Shared Ethernet Adapter, Shared Ethernet Adapter failover, and VLAN.

Virtual Ethernet technology facilitates IP-based communication between logical partitions on the same system by using software switch systems that are capable of virtual local area networks (VLANs). Using Shared Ethernet Adapter technology, logical partitions can communicate with other systems outside the hardware unit without assigning physical Ethernet slots to the logical partitions.

Host Ethernet adapter

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Note: HEA is not supported on POWER8 processor-based server.

Configuration of LHEA

Unlike most other types of I/O devices, you can never assign the HEA to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA, without having to go through an Ethernet bridge on another logical partition.

To connect a logical partition to an HEA, you must create a *logical Host Ethernet Adapter (LHEA)* for the logical partition. An LHEA appears to the operating system as if it were a physical Ethernet adapter, just as a virtual Ethernet adapter appears as if it were a physical Ethernet adapter. When you create an LHEA for a logical partition, you specify the resources that the logical partition can use on the actual physical HEA. Each logical partition can have one LHEA for each physical HEA on the managed system. Each LHEA can have one or more logical ports, and each logical port can connect to a physical port on the HEA.

You can create an LHEA for a logical partition by using either of the following methods:

- You can add the LHEA to a partition profile, shut down the logical partition, and reactivate the logical partition by using the partition profile with the LHEA.
- You can add the LHEA to a running logical partition by using dynamic partitioning for the following Linux logical partitions:

Table 19. Supported versions of Linux logical partitions

Linux logical partition	Supported versions
Red Hat Enterprise Linux	Version 4.6, or later Version 5.1, or later
SUSE Linux Enterprise Server	Version 10, or later Version 11, or later

When you activate a logical partition, the LHEAs in the partition profile are considered to be necessary resources. If the physical HEA resources that are necessary for the LHEAs are not available, the logical partition cannot be activated. However, when the logical partition is active, you can remove any LHEAs you want from the logical partition. For every active LHEA that you assign to an IBM® i logical partition, IBM i requires 40 MB of memory.

After you create an LHEA for a logical partition, a network device is created in the logical partition. This network device is named entX on AIX® logical partitions, CMNXX on IBM i logical partitions, and ethX on Linux logical partitions, where X represents sequentially assigned numbers. The user can then set up TCP/IP configuration like a physical Ethernet device to communicate with other logical partitions.

You can configure a logical partition so that it is the only logical partition that can access a physical port of an HEA by specifying promiscuous mode for an LHEA that is assigned to the logical partition. When an LHEA is in promiscuous mode, no other logical partitions can access the logical ports of the physical port that is associated with the LHEA that is in promiscuous mode. You might want to configure a logical partition to promiscuous mode in the following situations:

- If you want to connect more than 16 logical partitions to each other and to an external network through a physical port on an HEA, you can create a logical port on a Virtual I/O Server and configure an Ethernet bridge between the logical port and a virtual Ethernet adapter on a virtual LAN. This allows all logical partitions with virtual Ethernet adapters on the virtual LAN to communicate with the physical port through the Ethernet bridge. If you configure an Ethernet bridge between a logical port and a virtual Ethernet adapter, the physical port that is connected to the logical port must have the following properties:
  - The physical port must be configured so that the Virtual I/O Server is the promiscuous mode logical partition for the physical port.
  - The physical port can have only one logical port.
- You want the logical partition to have dedicated access to a physical port.
- You want to use tools such as *tcpdump* or *iptrace*.

A logical port can communicate with all other logical ports that are connected to the same physical port on the HEA. The physical port and its associated logical ports form a logical Ethernet network. Broadcast and multicast packets are distributed on this logical network as though it was a physical Ethernet network. You can connect up to 16 logical ports to a physical port by using this logical network. By extension, you can connect up to 16 logical partitions to each other and to an external network through this logical network. The actual number of logical ports that you can connect to a physical port depends upon the Multi-Core Scaling value of the physical port group. It also depends on the number of logical ports that are created for other physical ports within the physical port group. By default, the Multi-Core Scaling value of each physical port group is set to 4, which allows four logical ports to be connected to the physical ports in the physical port group. To allow up to 16 logical ports to be connected to the physical ports in the physical port group, you must change the Multi-Core Scaling value of the physical port group to 1 and restart the managed system.

You can set each logical port to restrict or allow packets that are tagged for specific VLANs. You can set a logical port to accept packets with any VLAN ID, or you can set a logical port to accept only the VLAN IDs that you specify. You can specify up to 20 individual VLAN IDs for each logical port.

The physical ports on an HEA are always configured on the managed system level. If you use an HMC to manage a system, you must use the HMC to configure the physical ports on any HEAs belonging to the managed system. Also, the physical port configuration applies to all logical partitions that use the physical port. (Some properties might require setup in the operating system as well. For example, the maximum packet size for a physical port on the HEA must be set on the managed system level by using the HMC. However, you must also set the maximum packet size for each logical port within the operating system.) By contrast, if a system is unpartitioned and is not managed by an HMC, you can configure the physical ports on an HEA within the operating system as if the physical ports were ports on a regular physical Ethernet adapter.

HEA hardware does not support half-duplex mode.

You can change the properties of a logical port on an LHEA by using dynamic partitioning to remove the logical port from the logical partition. You can also add the logical port back to the logical partition by using the changed properties. If the operating system of the logical partition does not support dynamic partitioning for LHEAs, and you want to change any logical port property other than the VLANs on which the logical port participates, you must set a partition profile for the logical partition so that the partition profile contains the wanted logical port properties, shut down the logical partition, and activate the logical partition by using the new or changed partition profile. If the operating system of the logical partition does not support dynamic partitioning for LHEAs, and you want to change the VLANs on which the logical port participates, you must remove the logical port from a partition profile that belongs to the logical partition, shut down and activate the logical partition by using the changed partition profile, add the logical port back to the partition profile by using the changed VLAN configuration, and shut down and activate the logical partition again by using the changed partition profile.

Internet Protocol version 6

Internet Protocol version 6 (IPv6) is the next generation of Internet Protocol and is gradually replacing the current internet standard, Internet Protocol version 4 (IPv4). The key IPv6 enhancement is the expansion of the IP address space 32 - 128 bits, providing unlimited, unique IP addresses.

IPv6 provides several advantages over IPv4, including expanded routing and addressing, routing simplification, header format simplification, improved traffic control, autoconfiguration, and security.

For more information about IPv6, see the following resources:

- AIX: Internet Protocol (IP) version 6
- IBM i: Internet Protocol version 6

Note: For more information about IPv6 on the Linux operating system, see the documentation for the Linux operating system.

Link aggregation or Etherchannel devices

A link aggregation, or Etherchannel device, is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters that are aggregated can then act as a single Ethernet device. Link aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` adapters can be aggregated to the `ent3` adapter. The system considers these aggregated adapters as one adapter, and all adapters in the link aggregation device are given the same hardware address. Therefore, they are treated by remote systems as if they were one adapter.

Link aggregation can provide increased redundancy because individual links might fail. The link aggregation device can automatically fail over to another adapter in the device to maintain connectivity. For example, if the `ent0` adapter fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. The `ent0` adapter automatically returns to service on the link aggregation device when it recovers.

You can configure a Shared Ethernet Adapter to use a link aggregation, or Etherchannel, device as the physical adapter.

Virtual Ethernet adapters

Virtual Ethernet adapters allow client logical partitions to send and receive network traffic without having a physical Ethernet adapter.

Virtual Ethernet adapters allow logical partitions within the same system to communicate without having to use physical Ethernet adapters. Within the system, virtual Ethernet adapters are connected to an IEEE 802.1Q virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VIDs. With VIDs, virtual Ethernet adapters can share a common logical network. The system transmits packets by copying the packet directly from the memory of the sender logical partition to the receive buffers of the receiver logical partition without any intermediate buffering of the packet.

You can use virtual Ethernet adapters without using the Virtual I/O Server, but the logical partitions cannot communicate with external systems. However, in this situation, you can use another device, called a Host Ethernet Adapter (or Integrated Virtual Ethernet) to facilitate communication between logical partitions on the system and external networks.

You can create virtual Ethernet adapters with the Hardware Management Console (HMC) and configure them using the Virtual I/O Server command-line interface. With the Virtual I/O Server Version 2.2, or later, you can add, remove, or modify the existing set of VLANs for a virtual Ethernet adapter that is assigned to an active partition on a POWER7, POWER8, or POWER9 processor-based servers by using the HMC. The server firmware level must be at least AH720_064+ for high end servers, AM720_064+ for midrange servers, and AL720_064+ for low end servers. The HMC must be at Version 7.7.2.0, with mandatory fix MH01235, or later, to perform this task.

Note: The AL720_064+ server firmware level is only supported on POWER7 processor-based servers, or later.

Consider using virtual Ethernet in the following situations:

- When the capacity or the bandwidth requirement of the individual logical partition is inconsistent with, or is less than, the total bandwidth of a physical Ethernet adapter. If logical partitions use the full bandwidth or capacity of a physical Ethernet adapter, use dedicated Ethernet adapters.
- When you need an Ethernet connection, but there is no slot available in which to install a dedicated adapter.

Virtual local area networks

Virtual local area networks (VLAN) allow the physical network to be logically segmented.

A VLAN is a method to logically segment a physical network so that layer 2 connectivity is restricted to members that belong to the same VLAN. This separation is achieved by tagging Ethernet packets with their VLAN membership information and then restricting delivery to members of that VLAN. VLAN is described by the IEEE 802.1Q standard.

The VLAN tag information is referred to as VLAN ID (VID). Ports on a switch are configured as being members of a VLAN designated by the VID for that port. The default VID for a port is referred to as the Port VID (PVID). The VID can be added to an Ethernet packet either by a VLAN-aware host, or by the switch in the case of VLAN-unaware hosts. Therefore, ports on an Ethernet switch must be configured with information that indicates whether the host connected is VLAN-aware.

For VLAN-unaware hosts, a port is set up as untagged and the switch tags all packets that enter through that port with the Port VLAN ID (PVID). The switch also untags all packets that exit that port before delivery to the VLAN unaware host. A port that is used to connect VLAN-unaware hosts is called an *untagged port*, and it can be a member of only a single VLAN identified by its PVID. Hosts that are VLAN-aware can insert and remove their own tags and can be members of more than one VLAN. These hosts are typically attached to ports that do not remove the tags before the packets are delivered to the host. However, it inserts the PVID tag when an untagged packet enters the port. A port allows only packets that are untagged or tagged with the tag of one of the VLANs that the port belongs to. These VLAN rules are in addition to the regular media access control (MAC) address-based forwarding rules followed by a switch. Therefore, a packet with a broadcast or multicast destination MAC is also delivered to member ports that belong to the VLAN that is identified by the tags in the packet. This mechanism ensures the logical separation of the physical network that is based on membership in a VLAN.

Shared Ethernet Adapters

With Shared Ethernet Adapters on the Virtual I/O Server logical partition, virtual Ethernet adapters on client logical partitions can send and receive outside network traffic.

A Shared Ethernet Adapter is a Virtual I/O Server component that bridges a physical Ethernet adapter and one or more virtual Ethernet adapters:

- The real adapter can be a physical Ethernet adapter, a Link Aggregation or Etherchannel device, a Logical Host Ethernet Adapter, or an SR-IOV logical port. The real adapter cannot be another Shared Ethernet Adapter or a VLAN pseudo-device.
- The virtual Ethernet adapter must be a virtual I/O Ethernet adapter. It cannot be any other type of device or adapter.
- All virtual Ethernet adapters in a Shared Ethernet Adapter must be members of the same virtual switch.

Using a Shared Ethernet Adapter, logical partitions on the virtual network can share access to the physical network and communicate with stand-alone servers and logical partitions on other systems. The Shared

Ethernet Adapter eliminates the need for each client logical partition to a dedicated physical adapter to connect to the external network.

A Shared Ethernet Adapter provides access by connecting the internal VLANs with the VLANs on the external switches. Using this connection, logical partitions can share the IP subnet with stand-alone systems and other external logical partitions. The Shared Ethernet Adapter forwards outbound packets that are received from a virtual Ethernet adapter to the external network and forwards inbound packets to the appropriate client logical partition over the virtual Ethernet link to that logical partition. The Shared Ethernet Adapter processes packets at layer 2, so the original MAC address and VLAN tags of the packet are visible to other systems on the physical network.

The Shared Ethernet Adapter has a bandwidth apportioning feature, also known as Virtual I/O Server quality of service (QoS). QoS allows the Virtual I/O Server to give a higher priority to some types of packets. In accordance with the IEEE 801.q specification, Virtual I/O Server administrators can instruct the Shared Ethernet Adapter to inspect bridged VLAN-tagged traffic for the VLAN priority field in the VLAN header. The 3-bit VLAN priority field allows each individual packet to be prioritized with a value in the range 0 - 7 to distinguish more important traffic from less important traffic. More important traffic is sent preferentially and uses more Virtual I/O Server bandwidth than less important traffic.

Note: When you use the trunk of the Virtual Ethernet Adapter on an HMC, only traffic on VLANs with specified VLAN IDs is delivered to the Virtual I/O Server with a VLAN tag. Consequently, to use this feature, the adapter must be configured with additional VLAN IDs when the trunk of the Virtual Ethernet Adapter is configured. Untagged traffic is always treated as though it belonged to the default priority class, that is, as if it had a priority value of 0.

Depending on the VLAN priority values found in the VLAN headers, packets are prioritized as follows.

- 1 (Least important)
- 2
- 0 (Default)
- 3
- 4
- 5
- 6
- 7 (Most important)

The Virtual I/O Server administrator can use QoS by setting the Shared Ethernet Adapter qos_mode attribute to either strict or loose mode. The default is disabled mode. The following definitions describe these modes:

disabled mode

This is the default mode. VLAN traffic is not inspected for the priority field. An example follows:

```
chdev -dev <SEA device name> -attr qos_mode=disabled
```

strict mode

More important traffic is sent preferentially over less important traffic. This mode provides better performance and more bandwidth to more important traffic; however, it can result in substantial delays for less important traffic. An example follows:

```
chdev -dev <SEA device name> -attr qos_mode=strict
```

loose mode

A cap is placed on each priority level so that after a number of bytes is sent for each priority level, the following level is serviced. This method ensures that all packets are eventually sent. More important traffic is given less bandwidth with this mode than with strict mode; however, the caps in loose mode are such that more bytes are sent for the more important traffic, so it still gets more bandwidth than less important traffic. An example follows:

```
chdev -dev <SEA device name> -attr qos_mode=loose
```

Notes:

- In either strict or loose mode, because the Shared Ethernet Adapter uses several threads to bridge traffic, it is still possible for less important traffic from one thread to be sent before more important traffic of another thread.
- The SR-IOV logical port that is created on VIOS as part of the dedicated virtual NIC configuration cannot be used as an SEA backing device.

For more information, see Managing virtual Network Interface Controllers.

GARP VLAN Registration Protocol

Shared Ethernet Adapters, in Virtual I/O Server Version 1.4 or later, support GARP VLAN Registration Protocol (GVRP), which is based on Generic Attribute Registration Protocol (GARP). GVRP allows for the dynamic registration of VLANs over networks, which can reduce the number of errors in the configuration of a large network. By propagating registration across the network through the transmission of Bridge Protocol Data Units (BPDUs), devices on the network have accurate knowledge of the bridged VLANs configured on the network.

When GVRP is enabled, communication travels one way, from the Shared Ethernet Adapter to the switch. The Shared Ethernet Adapter notifies the switch which VLANs can communicate with the network. The Shared Ethernet Adapter does not configure VLANs to communicate with the network based on information that is received from the switch. Rather, the configuration of VLANs that communicate with the network is statically determined by the virtual Ethernet adapter configuration settings.

Host Ethernet Adapter or Integrated Virtual Ethernet

A logical Host Ethernet Adapter (LHEA), which is sometimes referred to as Integrated Virtual Ethernet, is a physical adapter that you can use to configure virtual Ethernet. With Virtual I/O Server Version 1.4, or later, you can assign a logical host Ethernet port of an LHEA, as the real adapter of a Shared Ethernet Adapter. The logical host Ethernet port is associated with a physical port on the Host Ethernet Adapter. The Shared Ethernet Adapter uses the standard device driver interfaces provided by the Virtual I/O Server to communicate with the Host Ethernet Adapter.

To use a Shared Ethernet Adapter with a Host Ethernet Adapter, the following requirements must be met:

- The logical host Ethernet port must be the only port that is assigned to the physical port on the Host Ethernet Adapter. No other ports of the LHEA can be assigned to the physical port on the Host Ethernet Adapter.
- The LHEA on the Virtual I/O Server logical partition must be set to `promiscuous` mode. *Promiscuous mode* allows the LHEA (on the Virtual I/O Server) to receive all unicast, multicast, and broadcast network traffic from the physical network.

Suggestions

Consider using **Shared Ethernet Adapters** on the Virtual I/O Server in the following situations:

- When the capacity or the bandwidth requirement of the individual logical partition is inconsistent or is less than the total bandwidth of a physical Ethernet adapter. Logical partitions that use the full bandwidth or capacity of a physical Ethernet adapter must use dedicated Ethernet adapters.
- If you plan to migrate a client logical partition from one system to another.

Consider assigning a Shared Ethernet Adapter to a Logical Host Ethernet port when the number of Ethernet adapters that you need is more than the number of ports available on the LHEA, or you anticipate that your needs will grow beyond that number. If the number of Ethernet adapters that you need is fewer than or equal to the number of ports available on the LHEA, and you do not anticipate needing more ports in the future, you can use the ports of the LHEA for network connectivity rather than the Shared Ethernet Adapter.

Single root I/O virtualization

Single root I/O virtualization (SR-IOV) is a Peripheral component interconnect express (PCIe) standard architecture that defines extensions to PCIe specifications to enable multiple logical partitions running simultaneously within a system to share PCIe devices. The architecture defines virtual replicas of PCI functions known as virtual functions (VF). A Logical partition can connect directly to an SR-IOV adapter VF without going through a virtual intermediary (VI) such as a POWER Hypervisor or Virtual I/O Server. This ability provides for a low latency and lower CPU utilization alternative by avoiding a VI.

An SR-IOV capable adapter might be assigned to a logical partition in dedicated mode or enabled for shared mode. The management console provides an interface to enable SR-IOV shared mode. An SR-IOV capable adapter in shared mode is assigned to the POWER Hypervisor for management of the adapter and provisioning of adapter resources to logical partitions. The management console, along with the POWER Hypervisor, provides the ability to manage the adapter's physical Ethernet ports and logical ports. To connect a logical partition to an SR-IOV Ethernet adapter VF, create an SR-IOV Ethernet logical port for the logical partition. When you create an Ethernet logical port for a partition, select the adapter physical Ethernet port to connect to the logical partition and specify the resource requirements for the logical port. Each logical partition can have one or more logical ports from each SR-IOV adapter in shared mode. The number of logical ports for all configured logical partitions cannot exceed the adapter logical port limit.

To create an SR-IOV Ethernet logical port for a logical partition, use one of the following methods:

- Create an Ethernet logical port when you create a partition.
- Add an Ethernet logical port to a partition profile, shut down the logical partition, and reactivate the logical partition by using the partition profile.
- Add an Ethernet logical port to a running logical partition by using dynamic partitioning.

  Note: An SR-IOV adapter does not support Live Partition Mobility unless the VF is assigned to a shared Ethernet adapter.

When you activate a logical partition, the logical ports in the partition profile are considered to be a required resource. If the physical adapter resources required by the logical port are not available, the logical partition cannot be activated. However, logical ports can be removed dynamically from other logical partition to make the required resources available to the logical partition.

For an SR-IOV adapter in shared mode, the physical port switch mode can be configured in Virtual Ethernet Bridge (VEB) or Virtual Ethernet Port Aggregator (VEPA) mode. If the switch mode is configured in VEB mode, the traffic between the logical ports is not visible to the external switch. If the switch mode is configured in VEPA mode, the traffic between logical ports must be routed back to the physical port by the external switch. Before you enable the physical port switch in VEPA mode, ensure that the switch attached to the physical port is supported and is enabled for reflective relay.

When you create an Ethernet logical port, you can select a promiscuous permission to allow the logical port to be configured as a promiscuous logical port by the logical partition. A promiscuous logical port receives all unicast traffic with a destination address that does not match the address of one of the other logical ports configured for the same physical port. The number of logical ports with promiscuous permission configured for logical partitions, active or shutdown, on a physical port is limited to minimize potential performance impact due to increased processor usage associated with promiscuous logical ports. The management console indicates the number of logical ports on the physical port that are allowed to have a promiscuous permission setting.

When bridging between virtual Ethernet adapters and a physical Ethernet adapter, an SR-IOV Ethernet logical port might be used as the physical Ethernet adapter to access the outside network. When a logical port is configured as the physical Ethernet adapter for bridging, the logical port must have the promiscuous permission enabled. For example, if you create a logical port for a Virtual I/O Server logical partition and the intent is to use the logical port as the physical adapter for the shared Ethernet adapter, you must select the promiscuous permission for the logical port.

Configuration requirements

Consider the following configuration requirements when an Ethernet logical port is used as the physical Ethernet device for shared Ethernet adapter bridging:

- When there is a requirement to divert all network traffic to flow through an external switch, consider the following requirements:
  - The POWER Hypervisor virtual switch must be set to the VEPA switching mode and the SR-IOV Ethernet adapter physical port switch mode must also be set to the VEPA switching mode.
  - In addition, the logical port is the only logical port that is configured for the physical port.
- When you create an Ethernet logical port you can specify a capacity value. The capacity value specifies the required capacity of the logical port as a percentage of the capability of the physical port. The capacity value determines the amount of resources assigned to the logical port from the physical port. The assigned resources determine the minimum capability of the logical port. Physical port resources not used by other logical ports might be temporarily used by the logical port when the logical port exceeds its assigned resources to allow additional capability. System or network limitations can influence the amount of throughput a logical port can actually achieve. The maximum capacity that can be assigned to a logical port is 100%. The sum of the capacity values for all the configured logical ports on a physical port must be less than or equal to 100%. To minimize the configuration effort while adding additional logical ports, you might want to reserve physical port capacity for additional logical ports.
- When an Ethernet logical port is used as a physical adapter for bridging virtual Ethernet adapters, the parameter values such as the number of client virtual adapters and expected throughput must be considered when choosing a capacity value.
- The Ethernet logical ports allow the logical port to run diagnostics on the adapter and physical port. Select this permission only while running the diagnostics by using the logical port.

Related information

Adding a single root I/O virtualization logical port to a logical partition dynamically

Shutting down and restarting logical partitions

Creating logical partitions

Assigning a single root I/O virtualization logical port to a logical partition

Updating the SR-IOV adapter firmware

Shared memory

Shared memory is physical memory that is assigned to the shared memory pool and shared among multiple logical partitions. The *shared memory pool* is a defined collection of physical memory blocks that are managed as a single memory pool by the hypervisor. Logical partitions that you configure to use shared memory, share the memory in the pool with other shared memory partitions.

For example, you create a shared memory pool with 16 GB of physical memory. You then create three logical partitions, configure them to use shared memory, and activate the shared memory partitions. Each shared memory partition can use the 16 GB that are in the shared memory pool.

The hypervisor determines the amount of memory that is allocated from the shared memory pool to each shared memory partition based on the workload and memory configuration of each shared memory partition. When allocating the physical memory to the shared memory partitions, the hypervisor ensures that each shared memory partition can access only the memory that is allocated to the shared memory partition at any given time. A shared memory partition cannot access the physical memory that is allocated to another shared memory partition.

The amount of memory that you assign to the shared memory partitions can be greater than the amount of memory in the shared memory pool. For example, you can assign 12 GB to shared memory partition 1, 8 GB to shared memory partition 2, and 4 GB to shared memory partition 3. Together, the shared memory partitions use 24 GB of memory, but the shared memory pool has only 16 GB of memory. In this situation, the memory configuration is considered over committed.

Over committed memory configurations are possible because the hypervisor virtualizes and manages all of the memory for the shared memory partitions in the shared memory pool as follows:

1. When shared memory partitions are not actively using their memory pages, the hypervisor allocates those unused memory pages to shared memory partitions that currently need them. When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the

amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time. The hypervisor need not store any data in auxiliary storage.

2. When a shared memory partition requires more memory than the hypervisor can provide to it by allocating unused portions of the shared memory pool, the hypervisor stores some of the memory that belongs to a shared memory partition in the shared memory pool and stores the remainder of the memory that belongs to the shared memory partition in auxiliary storage. When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage. When the operating system attempts to access the data, the hypervisor might need to retrieve the data from auxiliary storage before the operating system can access it.

Because the memory that you assign to a shared memory partition might not always reside in the shared memory pool, the memory that you assign to a shared memory partition is *logical memory*. Logical memory is the address space assigned to a logical partition, that the operating system perceives as its main storage. For a shared memory partition, a subset of the logical memory is backed up by physical main storage (or physical memory from the shared memory pool) and the remaining logical memory is kept in auxiliary storage.

A Virtual I/O Server logical partition provides access to the auxiliary storage, or paging space devices, which are required for shared memory partitions in an over committed memory configuration. A *paging space device* is a physical or logical device that is used by a Virtual I/O Server to provide the paging space for a shared memory partition. The *paging space* is an area of nonvolatile storage that is used to hold portions of a shared memory partition's logical memory that does not reside in the shared memory pool. When the operating system that runs in a shared memory partition attempts to access data, and the data is located in the paging space device that is assigned to the shared memory partition, the hypervisor sends a request to a Virtual I/O Server to retrieve the data and write it to the shared memory pool so that the operating system can access it.

On systems that are managed by a Hardware Management Console (HMC), you can assign up to two Virtual I/O Server (VIOS) logical partitions to the shared memory pool at a time. When you assign two paging VIOS partitions to the shared memory pool, you can configure the paging space devices such that both paging VIOS partitions have access to the same paging space devices. When one paging VIOS partition becomes unavailable, the hypervisor sends a request to the other paging VIOS partition to retrieve the data on the paging space device.

You cannot configure paging VIOS partitions to use shared memory. Paging VIOS partitions do not use the memory in the shared memory pool. You assign paging VIOS partitions to the shared memory pool so that they can provide access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool.

Driven by workload demands from the shared memory partitions, the hypervisor manages over committed memory configurations by continually performing the following tasks:

- Allocating portions of physical memory from the shared memory pool to the shared memory partitions as needed.
- Requesting a paging VIOS partition to read and write data between the shared memory pool and the paging space devices as needed.

The ability to share memory among multiple logical partitions is known as the PowerVM Active Memory Sharing technology. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

Related reference

Configuration requirements for shared memory

Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

Related information

Paging space device

Paging VIOS partition

A Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*) provides access to the paging space devices for the logical partitions that are assigned to the shared memory pool (hereafter referred to as *shared memory partitions*).

When the operating system that runs in a shared memory partition attempts to access data, and the data is located in the paging space device that is assigned to the shared memory partition, the hypervisor sends a request to a paging VIOS partition to retrieve the data and write it to the shared memory pool so that the operating system can access it.

A paging VIOS partition is not a shared memory partition and does not use the memory in the shared memory pool. A paging VIOS partition provides access to the paging space devices for the shared memory partitions.

HMC

On systems that are managed by a Hardware Management Console (HMC), you can assign one or two paging VIOS partitions to the shared memory pool. When you assign a single paging VIOS partition to the shared memory pool, the paging VIOS partition provides access to all of the paging space devices for the shared memory partitions. The paging space devices can be located in physical storage in the server or on a storage area network (SAN). When you assign two paging VIOS partitions to the shared memory pool, you can configure each paging VIOS partition to access paging space devices in one of the following ways:

- You can configure each paging VIOS partition to access independent paging space devices. Paging space devices that are accessed by only one paging VIOS partition, or independent paging space devices, can be located in physical storage in the server or on a SAN.
- You can configure both paging VIOS partitions to access the same, or common, paging space devices. In this configuration, the paging VIOS partitions provide redundant access to paging space devices. When one paging VIOS partition becomes unavailable, the hypervisor sends a request to the other paging VIOS partition to retrieve the data on the paging space device. Common paging space devices must be located on a SAN to enable symmetrical access from both paging VIOS partitions.
- You can configure each paging VIOS partition to access some independent paging space devices and some common paging space devices.

If you configure the shared memory pool with two paging VIOS partitions, you can configure a shared memory partition to use either a single paging VIOS partition or redundant paging VIOS partitions. When you configure a shared memory partition to use redundant paging VIOS partitions, you assign a primary paging VIOS partition and a secondary paging VIOS partition to the shared memory partition. The hypervisor uses the primary paging VIOS partition to access the shared memory partition's paging space device. At this point, the primary paging VIOS partition is the current paging VIOS partition for the shared memory partition. The current paging VIOS partition is the paging VIOS partition that the hypervisor uses at any point in time to access data in the paging space device that is assigned to the shared memory partition. If the primary paging VIOS partition becomes unavailable, the hypervisor uses the secondary paging VIOS partition to access the shared memory partition's paging space device. At this point, the secondary paging VIOS partition becomes the current paging VIOS partition for the shared memory partition and continues as the current paging VIOS partition even after the primary paging VIOS partition becomes available again.

You do not need to assign the same primary and secondary paging VIOS partitions to all of the shared memory partitions. For example, you assign paging VIOS partition A and paging VIOS partition B to the shared memory pool. For one shared memory partition, you can assign paging VIOS partition A as the primary paging VIOS partition and paging VIOS partition B as the secondary paging VIOS partition. For a different shared memory partition, you can assign paging VIOS partition B as the primary paging VIOS partition and paging VIOS partition A as the secondary paging VIOS partition.

The following figure shows an example of a system with four shared memory partitions, two paging VIOS partitions, and four paging space devices.

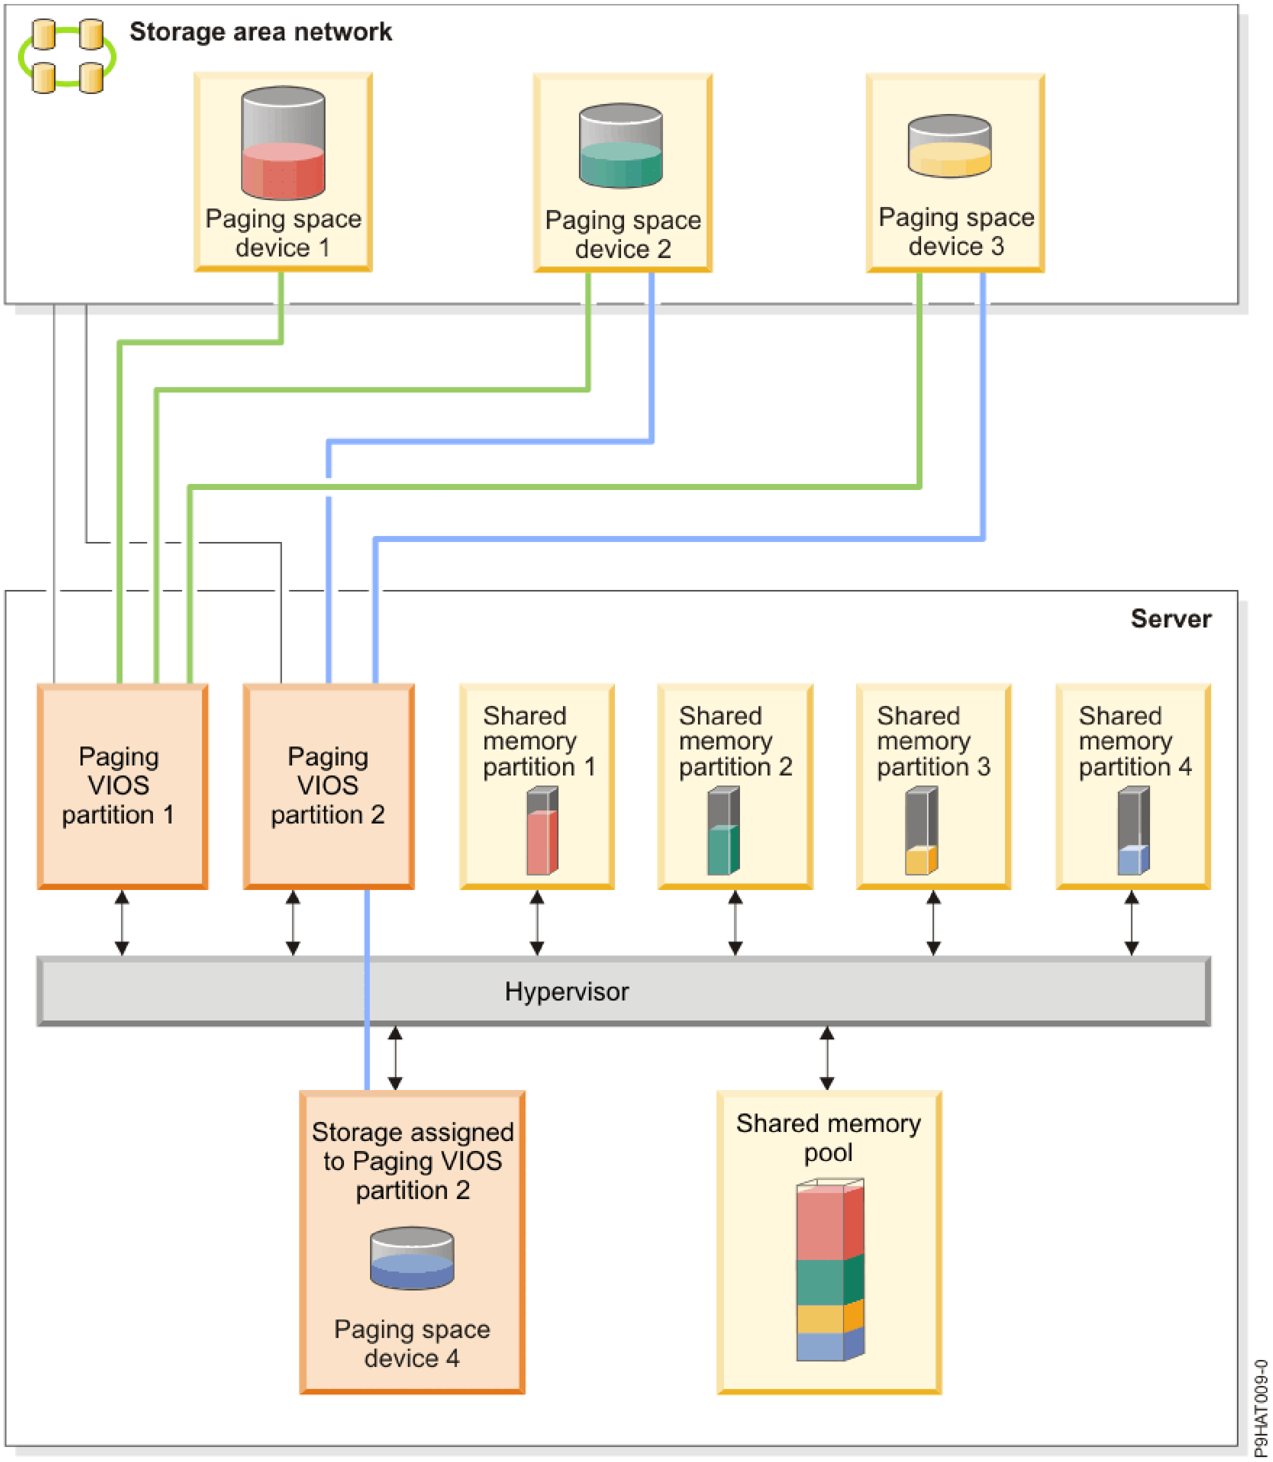

The example shows the configuration options for paging VIOS partitions and paging space devices as described in the following table.

Table 20. Examples of paging VIOS partition configurations

Configuration option	**Example**
The paging space device that is assigned to a shared memory partition is located in physical storage in the server and is accessed by a single paging VIOS partition.	Paging space device 4 provides the paging space for Shared memory partition 4. Shared memory partition 4 is assigned to use Paging VIOS partition 2 to access Paging space device 4. Paging space device 4 is located in physical storage in the server and is assigned to Paging VIOS partition 2. Paging VIOS partition 2 is the only paging VIOS partition that can access Paging space device 4 (This relationship is shown by the blue line that connects Paging VIOS partition 2 to Paging space device 4.).
The paging space device that is assigned to a shared memory partition is located on a SAN and is accessed by a single paging VIOS partition.	Paging space device 1 provides the paging space for Shared memory partition 1. Shared memory partition 1 is assigned to use Paging VIOS partition 1 to access Paging space device 1. Paging space device 1 is connected to the SAN. Paging VIOS partition 1 is also connected to the SAN and is the only paging VIOS partition that can access Paging space device 1 (This relationship is shown by the green line that connects Paging VIOS partition 1 to Paging space device 1.).

Table 20. Examples of paging VIOS partition configurations (continued)

Configuration option	Example
The paging space device that is assigned to a shared memory partition is located on a SAN and is accessed redundantly by two paging VIOS partitions.	Paging space device 2 provides the paging space for Shared memory partition 2. Paging space device 2 is connected to the SAN. Paging VIOS partition 1 and Paging VIOS partition 2 are also connected to the SAN and can both access Paging space device 2. (These relationships are shown by the green line that connects Paging VIOS partition 1 to Paging space device 2 and the blue line that connects Paging VIOS partition 2 to Paging space device 2.) Shared memory partition 2 is assigned to use redundant paging VIOS partitions to access Paging space device 2. Paging VIOS partition 1 is configured as the primary paging VIOS partition and Paging VIOS partition 2 is configured as the secondary paging VIOS partition. Similarly, Paging space device 3 provides the paging space for Shared memory partition 3. Paging space device 3 is connected to the SAN. Paging VIOS partition 1 and Paging VIOS partition 2 are also connected to the SAN and can both access Paging space device 3. (These relationships are shown by the green line that connects Paging VIOS partition 1 to Paging space device 3 and the blue line that connects Paging VIOS partition 2 to Paging space device 3.) Shared memory partition 3 is assigned to use redundant paging VIOS partitions to access Paging space device 3. Paging VIOS partition 2 is configured as the primary paging VIOS partition and Paging VIOS partition 1 is configured as the secondary paging VIOS partition. Because Paging VIOS partition 1 and Paging VIOS partition 2 both have access to Paging space device 2 and Paging space device 3, Paging space device 2 and Paging space device 3 are common paging space devices that are accessed redundantly by Paging VIOS partition 1 and Paging VIOS partition 2. If Paging VIOS partition 1 becomes unavailable and Shared memory partition 2 needs to access data on its paging space device, the hypervisor sends a request to Paging VIOS partition 2 to retrieve the data on Paging space device 2. Similarly, if Paging VIOS partition 2 becomes unavailable and Shared memory partition 3 needs to access the data on its paging space device, the hypervisor sends a request to Paging VIOS partition 1 to retrieve the data on Paging space device 3.

Table 20. Examples of paging VIOS partition configurations (continued)

Configuration option	Example
A paging VIOS partition accesses both independent and common paging space devices.	Paging space device 1 and Paging space device 4 are independent paging space devices because only one paging VIOS partition accesses each. Paging VIOS partition 1 accesses Paging space device 1, and Paging VIOS partition 2 accesses Paging space device 4. Paging space device 2 and paging space device 3 are common paging space devices because both paging VIOS partitions access each. (These relationships are shown by the green and blue lines that connect the paging VIOS partitions to the paging space devices.) Paging VIOS partition 1 accesses the independent paging space device Paging space device 1, and also accesses the common paging space devices Paging space device 2 and Paging space device 3. Paging VIOS partition 2 accesses the independent paging space device Paging space device 4 and also accesses the common paging space devices Paging space device 2 and Paging space device 3.

When a single paging VIOS partition is assigned to the shared memory pool, you must shut down the shared memory partitions before you shut down the paging VIOS partition so that the shared memory partitions are not suspended when they attempt to access their paging space devices. When two paging VIOS partitions are assigned to the shared memory pool and the shared memory partitions are configured to use redundant paging VIOS partitions, you do not need to shut down the shared memory partitions to shut down a paging VIOS partition. When one paging VIOS partition is shut down, the shared memory partitions use the other paging VIOS partition to access their paging space devices. For example, you can shut down a paging VIOS partition and install VIOS updates without shutting down the shared memory partitions.

You can configure multiple VIOS logical partitions to provide access to paging space devices. However, you can only assign up to two of those VIOS partitions to the shared memory pool at any given time.

After you configure the shared memory partitions, you can later change the redundancy configuration of the paging VIOS partitions for a shared memory partition by modifying the partition profile of the shared memory partition and restarting the shared memory partition with the modified partition profile:

- You can change which paging VIOS partitions are assigned to a shared memory partition as the primary and secondary paging VIOS partitions.
- You can change the number of paging VIOS partitions that are assigned to a shared memory partition.

Virtual I/O Server management

Learn about management tools for the Virtual I/O Server, such as the Virtual I/O Server command-line interface, and several Tivoli® products that can manage different aspects of the Virtual I/O Server.

Virtual I/O Server command-line interface

Learn about accessing and using the Virtual I/O Server command-line interface.

The Virtual I/O Server is configured and managed through a command-line interface. All aspects of Virtual I/O Server administration can be accomplished through the command-line interface, including the following:

- Device management (physical, virtual, logical volume manager (LVM))

- Network configuration
- Software installation and update
- Security
- User management
- Maintenance tasks

The first time you log in to the Virtual I/O Server, use the **padmin** user ID, which is the prime administrator user ID. You will be prompted for a new password.

Restricted shell

After logging in, you will be placed into a restricted Korn shell. The restricted Korn shell works in the same way as a standard Korn shell, except that you cannot perform the following:

- Change the current working directory
- Set the value of the **SHELL**, **ENV**, or **PATH** variables
- Specify the path name of the command that contains a forward slash (/)
- Redirect output of a command by using any of the following characters: >, >|, <>, >>

As a result of these restrictions, you cannot execute commands that are not accessible to your **PATH** variables. In addition, these restrictions prevent you from sending command output directly to a file. Instead, command output can be piped to the **`tee`** command.

After you log in, you can type `help` to get information about the supported commands. For example, to get help on the **`errlog`** command, type `help errlog`.

Execution mode

The Virtual I/O Server command-line interface functions similarly to a standard command-line interface. Commands are issued with appropriate accompanying flags and parameters. For example, to list all adapters, type the following:

```
lsdev -type adapter
```

In addition, scripts can be run within the Virtual I/O Server command-line interface environment.

In addition to the Virtual I/O Server command-line interface commands, the following standard shell commands are provided.

Table 21. Standard shell commands and their functions

Command	Function
`awk`	Matches patterns and performs actions on them.
`cat`	Concatenates or displays files.
`chmod`	Changes file modes.
`cp`	Copies files.
`date`	Displays the date and time.
`grep`	Searches a file for a pattern.
`ls`	Displays the contents of a directory.
`mkdir`	Makes a directory.
`man`	Displays manual entries for the Virtual I/O Server commands.
`more`	Displays the contents of files one screen at a time.

Table 21. Standard shell commands and their functions (continued)

Command	Function
rm	Removes files.
sed	Provides a stream editor.
stty	Sets, resets, and reports workstation operating parameters.
tee	Displays the output of a program and copies it to a file.
vi	Edits files with full screen display.
wc	Counts the number of lines, words, bytes, and characters in a file.
who	Identifies the users who are currently logged in.

As each command is executed, the user log and the global command log are updated.

The user log contains a list of each Virtual I/O Server command, including arguments, that a user has executed. One user log for each user in the system is created. This log is located in the home directory of the user and can be viewed by using either the **cat** or the **vi** commands.

The global command log is made up of all the Virtual I/O Server command-line interface commands executed by all users, including arguments, the date and time the command was executed, and from which user ID it was executed. The global command log is viewable only by the **padmin** user ID, and it can be viewed by using the **lsgcl** command. If the global command log exceeds 1 MB, the log is truncated to 250 KB to prevent the file system from reaching its capacity.

Remote script

Secure Shell (SSH) is shipped with the Virtual I/O Server. Hence, scripts and commands can run remotely after an exchange of SSH keys. To set up and run the commands remotely, perform the following steps:

1. From the command line on the remote system, type the **ssh** command and verify that the Virtual I/O Server has been added as a known host. If not, you must perform the following steps to exchange ssh keys.

   ```
   # ssh padmin@<vios> ioscli ioslevel
   padmin@<vios>'s password:
   2.1.2.0
   ```

   Where <vios> is either the Virtual I/O Server host name or its TCP/IP address.
2. Generate the public ssh key on the remote system.
3. Transfer the ssh key to the Virtual I/O Server. The transfer can be done by using File Transfer Protocol (FTP).
4. On the Virtual I/O Server, type the following command to copy the public key to the .ssh directory:

   ```
    $ cat id_rsa.pub >> .ssh/authorized_keys
   ```
5. From the command line on the remote system, type the same **ssh** command from step 1 to add the Virtual I/O Server as a known host. The command prompts the user for a password if it has not already been added as a known host.
6. From the command line on the remote system, type the same **ssh** command from step 1 to verify that the **ssh** command can run without requiring the user to enter a password.

Related information

Virtual I/O Server commands

IBM Tivoli software and the Virtual I/O Server

Learn about integrating the Virtual I/O Server into your Tivoli environment for IBM Tivoli Application Dependency Discovery Manager, IBM Tivoli Monitoring, IBM Tivoli Storage Manager, IBM Tivoli Usage and Accounting Manager, IBM Tivoli Identity Manager, and Tivoli Storage Productivity Center.

IBM Tivoli Application Dependency Discovery Manager

IBM Tivoli Application Dependency Discovery Manager discovers infrastructure elements found in the typical data center, including application software, hosts and operating environments (including the Virtual I/O Server), network components (such as routers, switches, load balancers, firewalls, and storage), and network services (such as LDAP, NFS, and DNS). Based on the data it collects, IBM Tivoli Application Dependency Discovery Manager automatically creates and maintains application infrastructure maps that include runtime dependencies, configuration values, and change history. With this information, you can determine the interdependency among business applications, software applications, and physical components to help you ensure and improve application availability in your environment. For example, you can do the following tasks:

- You can isolate configuration-related application problems.
- You can plan for application changes to minimize or eliminate unplanned disruptions.
- You can create a shared topological definition of applications for use by other management applications.
- You can determine the effect of a single configuration change on a business application or service.
- You can see what changes take place in the application environment and where.

IBM Tivoli Application Dependency Discovery Manager includes an agent-free discovery engine, which means that the Virtual I/O Server does not require that an agent or client be installed and configured to be discovered by IBM Tivoli Application Dependency Discovery Manager. Instead, IBM Tivoli Application Dependency Discovery Manager uses discovery sensors that rely on open and secure protocols and access mechanisms to discover the data center components.

IBM Tivoli Identity Manager

With IBM Tivoli Identity Manager, you can manage identities and users across several platforms, including AIX systems, Windows systems, Solaris systems, and so on. With IBM Tivoli Identity Manager 4.7 and later, you can also include Virtual I/O Server users. IBM Tivoli Identity Manager provides a Virtual I/O Server adapter that acts as an interface between the Virtual I/O Server and the IBM Tivoli Identity Manager Server. The adapter might not be located on the Virtual I/O Server and the IBM Tivoli Identity Manager Server manages access to the Virtual I/O Server by using your security system.

The adapter runs as a service, independent of whether a user is logged on to the IBM Tivoli Identity Manager Server. The adapter acts as a trusted virtual administrator on the Virtual I/O Server, performing tasks like the following:

- Creating a user ID to authorize access to the Virtual I/O Server.
- Modifying an existing user ID to access the Virtual I/O Server.
- Removing access from a user ID. This deletes the user ID from the Virtual I/O Server.
- Suspending a user account by temporarily deactivating access to the Virtual I/O Server.
- Restoring a user account by reactivating access to the Virtual I/O Server.
- Changing a user account password on the Virtual I/O Server.
- Reconciling the user information of all current users on the Virtual I/O Server.
- Reconciling the user information of a particular user account on the Virtual I/O Server by performing a lookup.

IBM Tivoli Monitoring

Virtual I/O Server V1.3.0.1 (fix pack 8.1), includes the IBM Tivoli Monitoring System Edition for IBM Power Systems. With Tivoli Monitoring System Edition for Power Systems, you can monitor the health and availability of multiple Power Systems (including the Virtual I/O Server) from the Tivoli Enterprise Portal. Tivoli Monitoring System Edition for Power Systems gathers data from the Virtual I/O Server, including data about physical volumes, logical volumes, storage pools, storage mappings, network mappings, real memory, processor resources, mounted file system sizes, and so on. From the Tivoli Enterprise Portal, you can view a graphical representation of the data, use predefined thresholds to alert you on key metrics, and resolve issues based on suggestions that are provided by the Expert Advice feature of Tivoli Monitoring.

IBM Tivoli Storage Manager

Virtual I/O Server 1.4 includes the IBM Tivoli Storage Manager client. With Tivoli Storage Manager, you can protect Virtual I/O Server data from failures and other errors by storing backup and disaster-recovery data in a hierarchy of auxiliary storage. Tivoli Storage Manager can help protect computers that run various different operating environments, including the Virtual I/O Server, on various different hardware, including Power Systems servers. If you configure the Tivoli Storage Manager client on the Virtual I/O Server, you can include the Virtual I/O Server in your standard backup framework.

IBM Tivoli Usage and Accounting Manager

Virtual I/O Server 1.4 includes the IBM Tivoli Usage and Accounting Manager agent on the Virtual I/O Server. Tivoli Usage and Accounting Manager helps you track, allocate, and invoice your IT costs by collecting, analyzing, and reporting on the actual resources used by entities such as cost centers, departments, and users. Tivoli Usage and Accounting Manager can gather data from multi-tiered data centers that include Windows, AIX, HP/UX Sun Solaris, Linux, IBM i, and VMware operating systems, and Virtual I/O Server appliance.

Tivoli Storage Productivity Center

With Virtual I/O Server 1.5.2, you can configure the TotalStorage™ agents on the Virtual I/O Server. Tivoli Storage Productivity Center is an integrated, storage infrastructure management suite that is designed to help simplify and automate the management of storage devices, storage networks, and capacity utilization of file systems and databases. When you install and configure the Tivoli Storage Productivity Center agents on the Virtual I/O Server, you can use the Tivoli Storage Productivity Center user interface to collect and view information about the Virtual I/O Server. You can then perform the following tasks by using the Tivoli Storage Productivity Center user interface:

1. Run a discovery job for the agents on the Virtual I/O Server.
2. Run probes, run scans, and ping jobs to collect storage information about the Virtual I/O Server.
3. Generate reports by using the Fabric Manager and the Data Manager to view the storage information gathered.
4. View the storage information gathered by using the topology Viewer.

Related tasks

Configuring the IBM Tivoli agents and clients on the Virtual I/O Server
You can configure and start the IBM Tivoli Monitoring agent, IBM Tivoli Usage and Accounting Manager, the IBM Tivoli Storage Manager client, and the Tivoli Storage Productivity Center agents.

Related information

IBM Tivoli Application Dependency Discovery Manager Information Center

IBM Tivoli Identity Manager

IBM Tivoli Monitoring version 6.2.1 documentation

IBM Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

IBM Tivoli Storage Manager

IBM Tivoli Usage and Accounting Manager Information Center

IBM TotalStorage Productivity Center Information Center

Virtual I/O Server rules management

Virtual I/O Server (VIOS) rules management provides capabilities to simplify VIOS device configuration and setup. It provides predefined default device settings based on the best practice values for VIOS. It also provides flexibility to manage and customize device settings.

You can collect, apply, and verify device settings in a VIOS runtime environment, by using VIOS rules management. It supports consistent device settings on multiple Virtual I/O Servers and updates, and also improves usability and ease of use of VIOS.

Rules file can be distributed to one or many VIOS partitions in a customer data center. This provides consistency between groups of VIOS partitions that use the same rules file. But VIOS rules file does not save nor preserve VIOS specific device instance information because device instance information might not apply to other Virtual I/O Servers.

Managing VIOS rules files

Virtual I/O Server (VIOS) rules management consists of two rules files. The *default rules file* contains the critical suggested device rules for VIOS best practice, and the *current rules file* captures the current VIOS system settings based on the default rules.

To deploy the suggested default device settings on a newly installed VIOS, run the `rules –o deploy` –d command and then restart the system. The default rules are contained in an XML profile, and you cannot modify the default rules.

You can customize rules on VIOS, by using the current rules. The initial current rules are captured from the system by using default rules as a template and then saving them in an XML profile. You can modify the current rules or add new rules. The new rules must be supported on the VIOS level. You can apply the changed current rules to VIOS, for currently discovered and newly discovered device types and instances.

You can use the **rules** command to manage VIOS rules files.

Viewing Virtual I/O Server rules

You can use the **–o list** option, with the **rules** command to view and list the contents of the default rules file, the current rules file, and the current system settings on the Virtual I/O Server. You can view the rules that are contained in a user-specified rules file by using the **–f** flag. The first column of the output describes a particular device in the *class/subclass/type* format. For example, `hdisk4` is described as `disk/fcp/osdisk`, where `disk` is the class, `hdisk4` might have the attribute *reserve_policy* with the value *single_path*.

Examples

1. To list the rules that are currently applied to the system, type the following command:

```
$ rules -o list -s
```

2. To list the rules in the current rules file, type the following command:

```
$ rules -o list
```

Deploying Virtual I/O Server rules

You can use the **rules** command with the **–o deploy** option to deploy rules. The **rules** command accepts the **–d** flag to deploy the default rules. Otherwise, the command uses the current rules on the VIOS. This command deploys the device type and then deploys the device instances by using the default or current rules. However, not all device instances on the system are deployed because of the VIOS specific configuration requirements. The new settings do not take effect until VIOS restarts.

Note: If your system does not have enough memory to accommodate the values for the attributes in the rules file, the rules are not deployed and a warning message is displayed.

To deploy the default rules on VIOS, type the following command:

```
$rules –o deploy -d
```

Capturing Virtual I/O Server rules

You can use the **–o capture** option, with the **rules** command to capture the current settings on VIOS. If the current rules file exists, it is used as the template to capture the latest system settings. If the VIOS has changed, this operation changes the current rules file.

To capture the current rules on VIOS, type the following command:

```
$rules –o capture
```

Importing Virtual I/O Server rules

You can use **-o import** option, with the **rules** command to import a user-specified rules file to VIOS. This operation might change the current rules. This operation merges the imported rules and the current rules. The user-specified rules precede the current rules during the merge operation. When a rule is not supported on the VIOS level, the import operation fails and displays a message to indicate that VIOS does not support a rule that is specified in the import file. You must remove the unsupported rule entries before you attempt the import operation again. A warning is displayed if the changed value is lower than the current default value in the AIX Object Data Manager (ODM). A lower value might impact performance or cause an LPM operation failure. If the **ioslevel** rule in the user-specified rules file is lower than the current rules, or if the ioslevel rule does not exist, the import operation stops. You can use the **-F** flag to force the import operation to continue and to ignore the **ioslevel** rule incompatibility.

To import the user rules file **user_rules.xml** to the current rules file on VIOS, type the following command:

```
$rules –o import –f user_rules.xml
```

Adding Virtual I/O Server rules

You can use the **-o add** option, with the **rules** command to add a rule entry to the VIOS current rules file or user-specified rules file, based on the *class/subclass/type* format or the device instance. If the rule that you are adding already exists on the VIOS, an error message is displayed to indicate that the rule already exists. The add operation might also fail if the VIOS level does not support a rule for the particular *class/subclass/type*, and if a template for the particular device does not exist. Currently, you can add only device rules. If the attribute value of the newly added rule is lower than the current Object Data Manager (ODM) default value, a warning message is displayed, but the operation is not stopped.

To add a rule for device type **cvdisk**, type the following command:

```
$ rules -o add -t disk/vscsi/cvdisk -a queue_depth=8
```

Modifying Virtual I/O Server rules

You can use the **-o modify** option, with the **rules** command to modify a rule from the current rules file or from the user specified file, based on the device type or device instance. If the rule that you want to modify does not exist in the current rules file, a message is displayed, prompting you to add the rule instead of modifying it. If the attribute value of the modified rule is lower than the current Object Data Manager (ODM) default value, a warning message is displayed, but the operation is not stopped.

To modify the *queue_depth* value of device type **cvdisk**, type the following command:

```
$ rules -o modify -t disk/vscsi/cvdisk -a queue_depth=16
```

Deleting Virtual I/O Server rules

You can use the **-o delete** option along with the **rules** command to delete a rule from the current rules file, or from the user-specified file, based on the device type or device instance.

Note: You cannot delete the rules that are taken from the current rules file and defined in the default rules file. The current rules file is used as a default template for capturing the rules, listing the rules, and other operations. After a rule is removed from the current rules file, that rule cannot be accessed for any rules operations that use the current rules file as the template.

To delete the rule for the *queue_depth* value of device type **cvdisk**, type the following command:

```
$ rules -o delete -t disk/vscsi/cvdisk -a queue_depth
```

Identifying mismatched rules on the Virtual I/O Server

You can use the **-o diff** operation to find the mismatched list of devices and attributes between the current rules file and the current VIOS settings, or between the default rules file and current rules file, or between the current VIOS setting and the default rules file. You can also detect the mismatched list between a rules file by specifying the **-f** flag with the current rules file, default rules file, or current system settings. If you use the **–n** flag, a count of all the mismatched list of devices and attributes is displayed.

Examples

1. To see the difference between the current rules file and rules that are currently applied on the system, type the following command:

   ```
   $ rules -o diff -s
   ```

2. To see the difference between the current rules file and the default rules file, type the following command:

   ```
   $ rules -o diff -d
   ```

Managing VIOS updates by using rules files

Virtual I/O Server (VIOS) updates might include updates to support new devices that might introduce new rules. VIOS replaces the default rules on the VIOS with the default rules file in the update media.

VIOS ships only one default rules file in each release. The default rules file contains accumulative changes for devices and attributes in successive updates.

The default rules file cannot be changed. However, if required, you can use it to set the system settings to factory default settings. However, if a setting does not exist in the default rules file, that setting does not reset.

When VIOS is upgraded from a level that does not support VIOS rules, the default rules file is copied to the current rules file. When VIOS is updated from a level that supports VIOS rules, the default rules file and the new device rules are merged into the current rules file, without overwriting the current rules. The current rules always precede the default rules. This ensures that the saved previous system settings remain unchanged.

After the update process completes, the current rules file can be used to restore the previous system configuration settings and you can apply new device rules to overwrite the existing rules.

If a mismatch is identified between default rules file and current rules file, a notification to call the **rulescfgset** command to apply updates is displayed. The new device rules are not applied until you run the **rulescfgset** command and type *yes* to confirm the deploy operation. The new device updates take effect after the VIOS reboots.

The notification can be disabled, by running the following command: `chdev –l viosrules0 –a motd=no`.

Managing EMC devices by using rules files

Virtual I/O Server (VIOS) provides the framework to manage EMC device configuration setup.

When the EMC software is installed, rules management merges the specific EMC default rules file into VIOS default rules file, and then it merges the default rules file to the current rules file. The current rules file precedes the default rules file.

If a mismatch is identified between system settings and current rules file, a notification to call the **rulescfgset** command to apply new EMC rules file is displayed. The EMC devices are not applied unless you run the **rulescfgset** command and type *yes* to confirm the deploy operation. The new EMC settings take effect after the VIOS reboots.

The notification can be disabled by running the following command: `chdev –l viosrules0 –a motd=no`.

When the EMC software is uninstalled, the specific default rules of the EMC device are removed from VIOS default rules and current rules files.

Distributing rules files to multiple VIOS partitions

To distribute the rules file to multiple VIOS partitions, complete the following steps:

1. Capture the current rules file from a source VIOS that contains the necessary configurations, by typing the following command:

   ```
   rules -o capture
   ```

2. Copy the current rules file **/home/padmin/rules/vios_current_rules.xml** from the source VIOS to the target Virtual I/O Servers.
3. Merge the current rules file from the source VIOS to the current rules file on the target Virtual I/O Servers, by typing the following command:

   ```
   rules –o import –f <curren_rules_file_from_source_vios>
   ```

4. Deploy the merged current rules on the target Virtual I/O Servers, by typing the following command:

   ```
   rules –o deploy
   ```

5. Restart the target Virtual I/O Servers as `padmin`, by typing the following command:

   ```
   shutdown -restart
   ```

Scenarios: Configuring the Virtual I/O Server

The following scenarios show examples of networking configurations for the Virtual I/O Server logical partition and the client logical partitions. Use the following scenarios and configuration examples to understand more about the Virtual I/O Server and its components.

Scenario: Configuring a Virtual I/O Server without VLAN tagging

Use this scenario to help you become familiar with creating a network without VLAN tagging.

About this task

Situation

You are the system administrator responsible for planning and configuring the network in an environment with the Virtual I/O Server running. You want to configure a single logical subnet on the system that communicates with the switch.

Objective

The objective of this scenario is to configure the network where only Port Virtual LAN ID (PVID) is used, the packets are not tagged, and a single internal network is connected to a switch. There are no virtual local area networks (VLAN) tagged ports set up on the Ethernet switch, and all virtual Ethernet adapters are defined by using a single default PVID and no additional VLAN IDs (VIDs).

Prerequisites and assumptions

- The Hardware Management Console (HMC) was set up. For more information about Installing and configuring the HMC, see Installing and configuring the Hardware Management Console.
- You understand the partitioning concepts as described in the Logical partitioning. For more information about Logical partitioning, see Logical partitioning.
- The Virtual I/O Server logical partition has been created and the Virtual I/O Server has been installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 90.
- You have created the remaining logical partitions that you want added to the network configuration.
- You have an Ethernet switch and a router ready to add to the configuration.
- You have IP addresses for all logical partitions and systems that are to be added to the configuration.

Configuration steps

The following figure shows the configuration that is to be completed during this scenario.

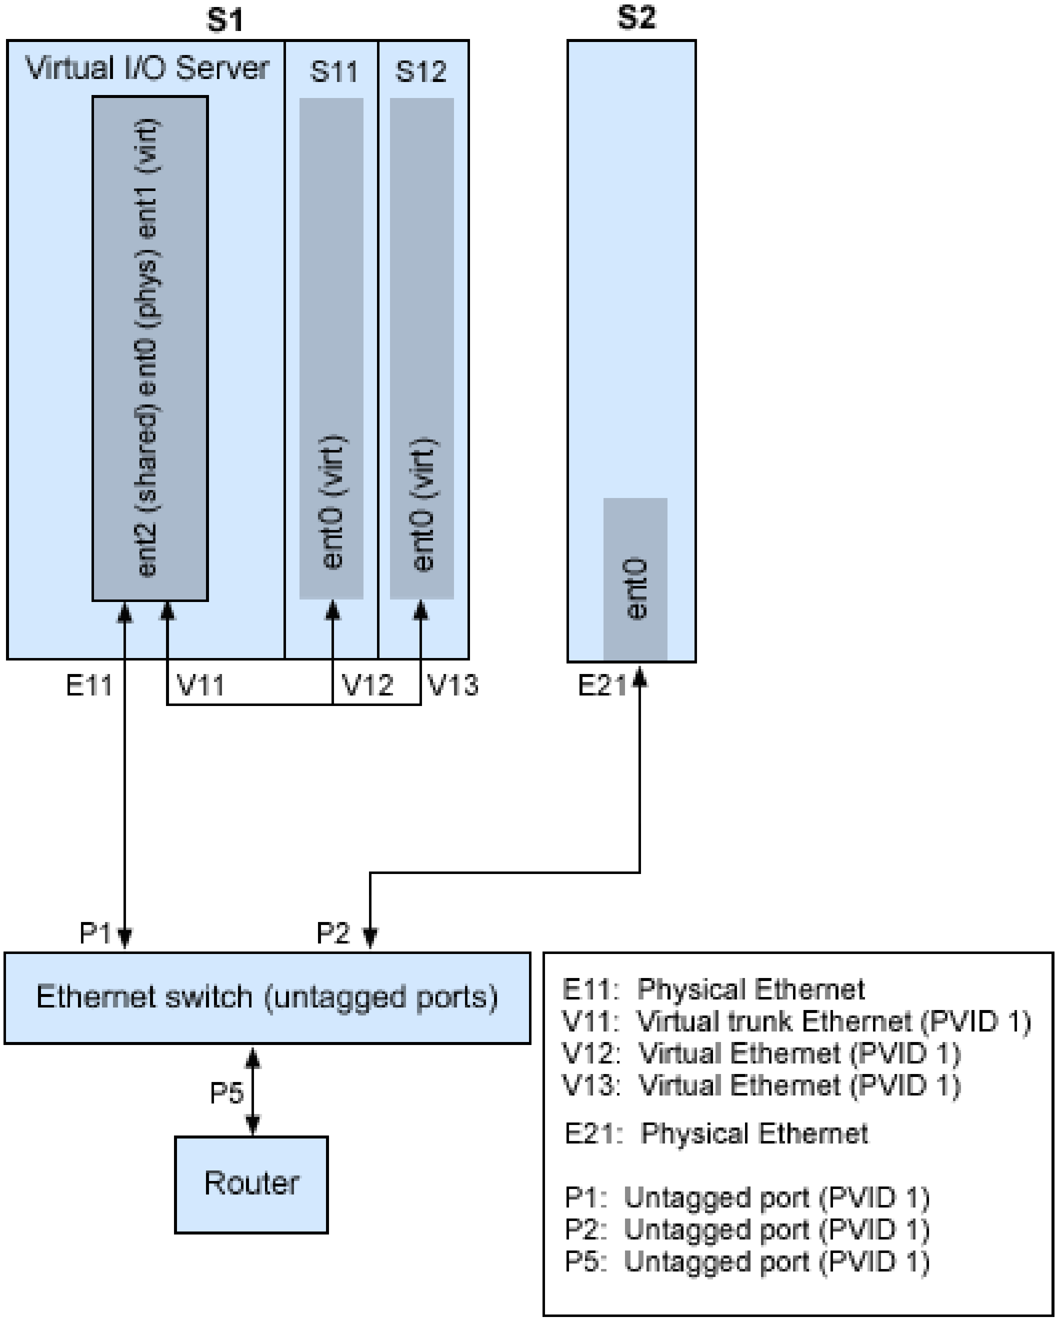
S1
S2
Virtual I/O Server
S11
S12
ent2 (shared) ent0 (phys) ent1 (virt)
ent0 (virt)
ent0 (virt)
ent0
E11
V11
V12
V13
E21
P1
P2
Ethernet switch (untagged ports)
P5
Router
E11: Physical Ethernet
V11: Virtual trunk Ethernet (PVID 1)
V12: Virtual Ethernet (PVID 1)
V13: Virtual Ethernet (PVID 1)
E21: Physical Ethernet
P1: Untagged port (PVID 1)
P2: Untagged port (PVID 1)
P5: Untagged port (PVID 1)

Using the preceding figure as a guide, follow these steps:

Procedure

1. Set up an Ethernet switch with untagged ports. Alternatively, you can use an Ethernet switch that does not use VLAN.
2. For system S1, use the HMC to create a virtual Ethernet adapter V11 for the Virtual I/O Server with the **Use this adapter for Ethernet bridging** trunk setting, with PVID set to 1, and no additional VIDs.
3. For system S1, use the HMC to create virtual Ethernet adapters V12 and V13 for logical partitions S11 and S12, with PVID set to 1 and no additional VIDs.
4. For system S1, use the HMC to assign physical Ethernet adapter E11 to the Virtual I/O Server and connect the adapter to the Ethernet switch port P1.

5. On the Virtual I/O Server, set up a shared Ethernet adapter (SEA) ent2 with the physical adapter ent0 and virtual adapter ent1 by using the `mkvdev -sea ent0 -vadapter ent1 -default ent1 -defaultid 1` command.
6. Start the logical partitions. The process recognizes the virtual devices that were created in Step 1.
7. Configure IP addresses for S11 (en0), S12 (en0), and S2 (en0), so that they all belong to the same subnet with the router connected to Ethernet switch port P5.

Results

An en2 SEA on the Virtual I/O Server logical partition can also be configured by using the IP addresses on the same subnet. This is required only for network connectivity to the Virtual I/O Server logical partition.

Scenario: Configuring a Virtual I/O Server by using VLAN tagging

Use this scenario to help you become familiar with creating a network by using VLAN tagging.

About this task

Situation

You are the system administrator responsible for planning and configuring the network in an environment with the Virtual I/O Server running. You would like to configure the network so that two logical subnets exist, with some logical partitions on each subnet.

Objective

The objective of this scenario is to configure multiple networks to share a single physical Ethernet adapter. Systems on the same subnet are required to be on the same VLAN, and therefore they have the same VLAN ID, which allows communication without having to go through the router. The separation in the subnets is achieved by ensuring that the systems on the two subnets have different VLAN IDs.

Prerequisites and assumptions

- The Hardware Management Console (HMC) is set up. For more information about installing and configuring the HMC, see Installing and configuring the Hardware Management Console.
- You understand the logical partitioning concepts. For more information, see Logical partitioning.
- The Virtual I/O Server logical partition is created and the Virtual I/O Server is installed. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 90.
- You created the remaining AIX, Linux or IBM i logical partitions that you want added to the network configuration. For more information, see Creating logical partitions. (VLAN tagging is supported in IBM i logical partitions Version 7.2, or later.)
- You have an Ethernet switch and a router ready to add to the configuration.
- You have IP addresses for all logical partitions and systems that are to be added to the configuration.

Configuration steps

The following figure shows the configuration that is to be completed during this scenario.

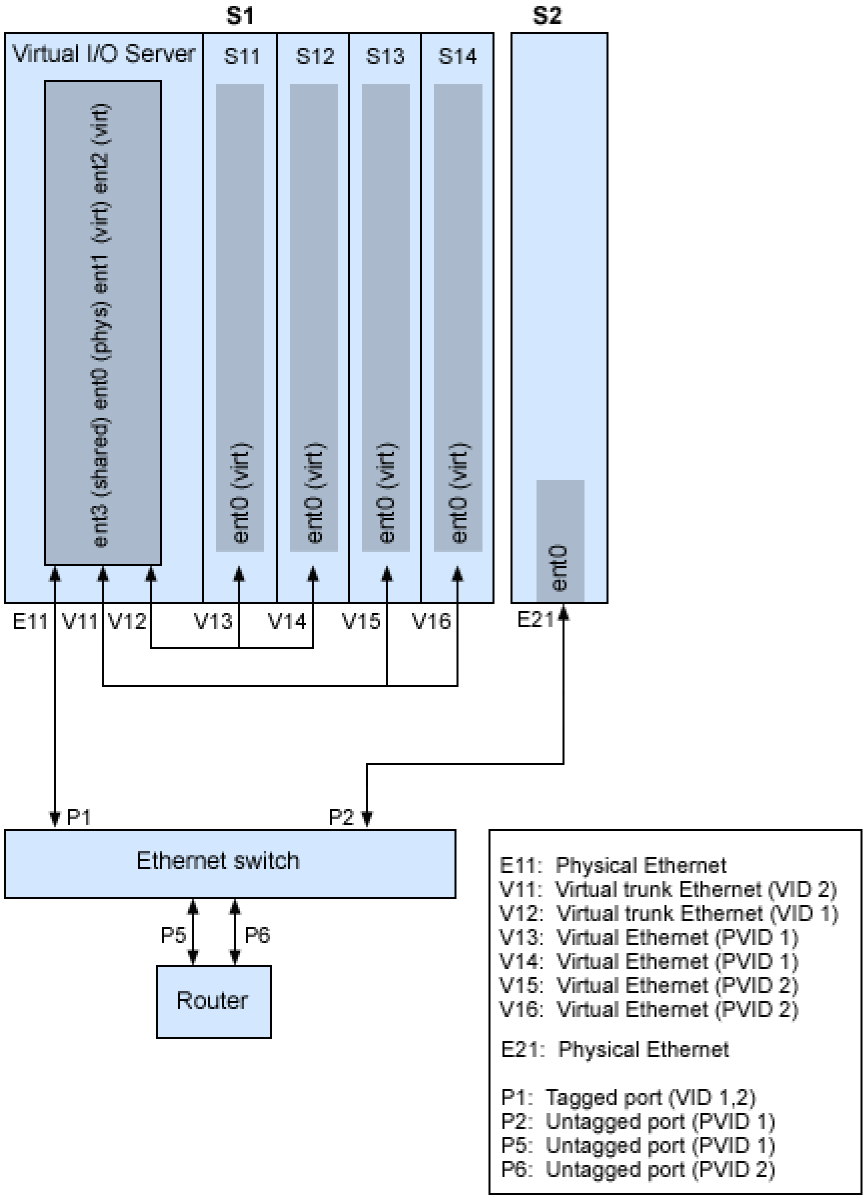
S1
S2
Virtual I/O Server
S11
S12
S13
S14
ent3 (shared) ent0 (phys) ent1 (virt) ent2 (virt)
ent0 (virt)
ent0 (virt)
ent0 (virt)
ent0 (virt)
ent0
E11
V11
V12
V13
V14
V15
V16
E21
P1
P2
Ethernet switch
P5
P6
Router
E11: Physical Ethernet
V11: Virtual trunk Ethernet (VID 2)
V12: Virtual trunk Ethernet (VID 1)
V13: Virtual Ethernet (PVID 1)
V14: Virtual Ethernet (PVID 1)
V15: Virtual Ethernet (PVID 2)
V16: Virtual Ethernet (PVID 2)
E21: Physical Ethernet
P1: Tagged port (VID 1,2)
P2: Untagged port (PVID 1)
P5: Untagged port (PVID 1)
P6: Untagged port (PVID 2)

Using the preceding figure as a guide, follow these steps.

Procedure

1. Set up the Ethernet switch ports as follows:
   - P1: Tagged port (VID 1, 2)
   - P2: Untagged port (PVID 1)
   - P5: Untagged port (PVID 1)
   - P6: Untagged port (PVID 2)

   For instructions about configuring the ports, see the documentation for your switch.
2. For system S1, use the HMC to create virtual Ethernet adapters for the Virtual I/O Server:
   - Create virtual Ethernet adapter V11 for the Virtual I/O Server with the trunk setting selected and VID set to 2. Specify an unused PVID value. This value is required, even though it is not used.
   - Create virtual Ethernet adapter V12 for the Virtual I/O Server with the trunk setting selected and VID set to 1. Specify an unused PVID value. This value is required, even though it is not used.
3. For system S1, use the HMC to create virtual Ethernet adapters for other logical partitions:

- Create virtual adapters V13 and V14 for logical partitions S11 and S12, with PVID set to 2 and no additional VIDs.
- Create virtual adapters V15 and V16 for logical partitions S13 and S14, with PVID set to 1 and no additional VIDs.

4. For system S1, use the HMC to assign the physical Ethernet adapter (E11) to the Virtual I/O Server and to connect the adapter to the Ethernet switch port P1.
5. Using the Virtual I/O Server command-line interface, set up a Shared Ethernet Adapter ent3 with the physical adapter ent0 and virtual adapters ent1 and ent2.
6. Configure IP addresses as follows:
   - S13 (ent0), S14 (ent0), and S2 (ent0) belong to VLAN 1 and are on the same subnet. The router is connected to Ethernet switch port P5.
   - S11 (ent0) and S12 (ent0) belong to VLAN 2 and are on the same subnet. The router is connected to Ethernet switch port P6.

Results

You can configure the Shared Ethernet Adapter on the Virtual I/O Server logical partition with an IP address. This is required only for network connectivity to the Virtual I/O Server.

As the tagged VLAN network is being used, you must define additional VLAN devices over the **Shared Ethernet Adapters** before configuring IP addresses.

Scenario: Configuring Shared Ethernet Adapter failover

Use this scenario to help you to configure primary and backup **Shared Ethernet Adapters** in the Virtual I/O Server logical partitions.

About this task

Situation

You are the system administrator responsible for planning and configuring the network in an environment with the Virtual I/O Server running. You want to provide higher network availability to the client logical partition on the system. This can be accomplished by configuring a backup Shared Ethernet Adapter in a different Virtual I/O Server logical partition.

Objective

The objective of this scenario is to configure primary and backup **Shared Ethernet Adapters** in the Virtual I/O Server logical partitions so that network connectivity in the client logical partitions will not be lost in the case of adapter failure.

Prerequisites and assumptions

- The Hardware Management Console (HMC) was set up. For more information about Installing and configuring the HMC, see Installing and configuring the Hardware Management Console.
- You understand the partitioning concepts as described in the Logical partitioning. For more information about Logical partitioning, see Logical partitioning.
- Two separate Virtual I/O Server logical partitions have been created and the Virtual I/O Server has been installed in each logical partition. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 90.
- You must have only one pair of Shared Ethernet Adapters per vSwitch/PVID, the VLAN IDs on this Shared Ethernet Adapter pair must not be on any other Shared Ethernet Adapter on the vSwitch, and you must understand what Shared Ethernet Adapter failover is and how it works. See “Shared Ethernet Adapter failover” on page 83.
- You have created the remaining logical partitions that you want added to the network configuration.
- EachVirtual I/O Server logical partition has an available physical Ethernet adapter assigned to it.

- You have IP addresses for all logical partitions and systems that will be added to the configuration.

The following image depicts a configuration where the Shared Ethernet Adapter failover feature is set up. The client logical partitions H1 and H2 are accessing the physical network using the **Shared Ethernet Adapters**, which are the primary adapters. The virtual Ethernet adapters used in the shared Ethernet setup are configured with the same VLAN membership information (PVID, VID), but have different priorities. A dedicated virtual network forms the control channel and is required to facilitate communication between the primary and backup shared Ethernet device.

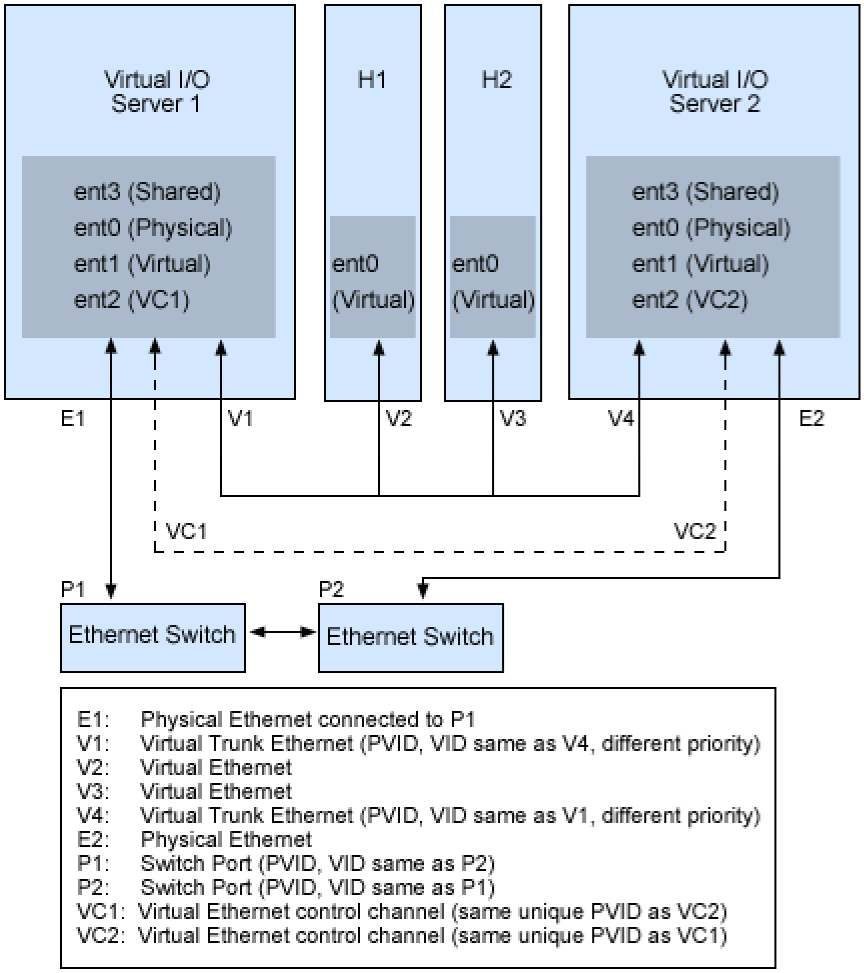
Virtual I/O Server 1
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
ent2 (VC1)
H1
ent0 (Virtual)
H2
ent0 (Virtual)
Virtual I/O Server 2
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
ent2 (VC2)
E1
V1
V2
V3
V4
E2
VC1
VC2
P1
P2
Ethernet Switch
Ethernet Switch
E1: Physical Ethernet connected to P1
V1: Virtual Trunk Ethernet (PVID, VID same as V4, different priority)
V2: Virtual Ethernet
V3: Virtual Ethernet
V4: Virtual Trunk Ethernet (PVID, VID same as V1, different priority)
E2: Physical Ethernet
P1: Switch Port (PVID, VID same as P2)
P2: Switch Port (PVID, VID same as P1)
VC1: Virtual Ethernet control channel (same unique PVID as VC2)
VC2: Virtual Ethernet control channel (same unique PVID as VC1)

Using the preceding figure as a guide, follow these steps:

Procedure

1. On the HMC, create the virtual Ethernet adapters following these guidelines:
   - Configure the virtual adapters to be used for data as trunk adapters by selecting the trunk setting.
   - Assign different prioritization values (valid values are 1-15) to each virtual adapter.
   - Configure another virtual Ethernet to be used for the control channel by giving it a unique PVID value. Make sure you use the same PVID when creating this virtual Ethernet for both Virtual I/O Server logical partitions.
2. Using the Virtual I/O Server command line, run the following command to configure the Shared Ethernet Adapter. Run this command on both Virtual I/O Server logical partitions involved in the configuration:

```
mkvdev -sea physical_adapter -vadapter virtual_adapter -default
virtual_adapter\
-defaultid PVID_of_virtual_adapter -attr ha_mode=auto
ctl_chan=control_channel_adapter
```

For example, in this scenario, run the following command on both Virtual I/O Server logical partitions:

```
mkvdev -sea ent0 -vadapter ent1 -default ent1 -defaultid 60 -attr ha_mode=auto
ctl_chan=ent2
```

Scenario: Configuring Shared Ethernet Adapter failover with load sharing

Use this scenario to help you to configure primary and backup **Shared Ethernet Adapters** for load sharing in the Virtual I/O Server (VIOS) logical partitions.

About this task

Situation

You are the system administrator responsible for planning and configuring the network in an environment with the VIOS running. You want to provide load sharing in addition to Shared Ethernet Adapter failover to improve the bandwidth of the VIOS logical partition without impact to higher network availability.

Objective

The objective of this scenario is to configure primary and backup **Shared Ethernet Adapters** for load sharing so that you can use both the **Shared Ethernet Adapters** by sharing the bridging workload between them.

Prerequisites and assumptions

- The Hardware Management Console (HMC) was set up. For more information about Installing and configuring the HMC, see Installing and configuring the Hardware Management Console.
- You understand the partitioning concepts as described in the Logical partitioning. For more information about Logical partitioning, see Logical partitioning.
- You have configured primary and backup **Shared Ethernet Adapters** in the VIOS logical partitions. See “Scenario: Configuring Shared Ethernet Adapter failover” on page 59.
- You understand what Shared Ethernet Adapter load sharing is and how it works. See “Shared Ethernet adapters for load sharing” on page 84.
- The VIOS must be at Version 2.2.1.0, or later.
- The VIOS servers with the primary and backup Shared Ethernet Adapter support load sharing.
- Two or more trunk adapters are configured for the primary and backup Shared Ethernet Adapter pair.
- The virtual local area network (VLAN) definitions of the trunk adapters are identical between the primary and backup Shared Ethernet Adapter pair.

Note: Enable load sharing mode on the primary Shared Ethernet Adapter (the Shared Ethernet Adapter with higher priority) before you enable load sharing mode on the backup Shared Ethernet Adapter (the Shared Ethernet Adapter with lesser priority).

To configure **Shared Ethernet Adapters** for load sharing, use the VIOS command line and run the following command. Run this command on both **Shared Ethernet Adapters**.

```
mkvdev -sea physical_adapter -vadapter virtual_adapter1, virtual_adapter2 -default
virtual_adapter1\
-defaultid PVID_of_virtual_adapter1 -attr ha_mode=sharing
ctl_chan=control_channel_adapter
```

For example, in this scenario, run the following command on both **Shared Ethernet Adapters**:

```
mkvdev -sea ent0 -vadapter ent1,ent2 -default ent1 -defaultid 60 -attr ha_mode=sharing
ctl_chan=ent3
```

What to do next

You can restart load sharing by using the **chdev** command on the backup Shared Ethernet Adapter. To restart load sharing, ensure that the **ha_mode** attribute is set to sharing on both the primary and backup Shared Ethernet Adapter. By using the VIOS command line, run the chdev command on the backup Shared Ethernet Adapter. If the load sharing criteria are met, load sharing restarts.

Scenario: Configuring Shared Ethernet Adapter failover without using a dedicated control channel adapter

Use this scenario to help you to configure Shared Ethernet Adapter failover in the Virtual I/O Server (VIOS) logical partitions without specifying the **Control Channel** attribute.

About this task

Situation

You are the system administrator responsible for planning and configuring the network in an environment with the VIOS running. You want to provide higher network availability to the client logical partition on the system. However, you do not want to use dedicated resources, such as a virtual Ethernet adapter and a virtual LAN that are required for the control channel adapter. This can be accomplished by configuring a Shared Ethernet Adapter in high availability mode in a VIOS logical partition without a dedicated control channel adapter.

Objective

The objective of this scenario is to configure a Shared Ethernet Adapter in high availability mode in the VIOS logical partitions without specifying the **Control Channel** attribute. This avoids the requirement of a dedicated virtual Ethernet adapter and a dedicated virtual LAN for the control channel adapter while you configure the Shared Ethernet Adapter in high availability mode.

Prerequisites and assumptions

- The Hardware Management Console (HMC) was set up. For more information about Installing and configuring the HMC, see Installing and configuring the Hardware Management Console.
- You must understand the partitioning concepts as described in the Logical partitioning. For more information about Logical partitioning, see Logical partitioning.
- You must have only one pair of Shared Ethernet Adapters per vSwitch/PVID, the VLAN IDs on this Shared Ethernet Adapter pair must not be on any other Shared Ethernet Adapter on the vSwitch, and you must understand what Shared Ethernet Adapter failover is and how it works. See "Shared Ethernet Adapter failover" on page 83.
- The Power Hypervisor must be at Version 780, or later.
- The VIOS must be at Version 2.2.3.0, or later.

Note: Even though the Power Hypervisor is at Version 780, configuring Shared Ethernet Adapter failover in the VIOS logical partitions, without specifying the **Control Channel** attribute is not supported on some of the servers, such as MMB servers and MHB servers.

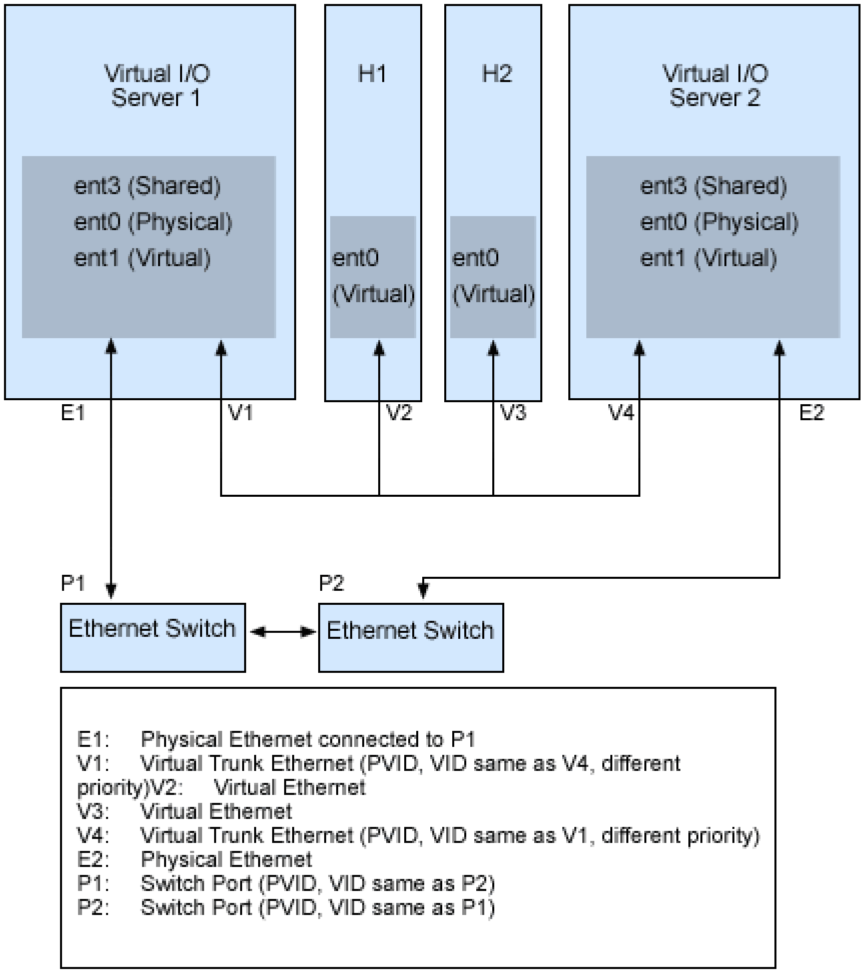
Virtual I/O Server 1
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
H1
ent0 (Virtual)
H2
ent0 (Virtual)
Virtual I/O Server 2
ent3 (Shared)
ent0 (Physical)
ent1 (Virtual)
E1
V1
V2
V3
V4
E2
P1
P2
Ethernet Switch
Ethernet Switch
E1: Physical Ethernet connected to P1
V1: Virtual Trunk Ethernet (PVID, VID same as V4, different priority)V2: Virtual Ethernet
V3: Virtual Ethernet
V4: Virtual Trunk Ethernet (PVID, VID same as V1, different priority)
E2: Physical Ethernet
P1: Switch Port (PVID, VID same as P2)
P2: Switch Port (PVID, VID same as P1)

In this configuration, the default adapter of the Shared Ethernet Adapter that is illustrated as V1 in the figure, is used as control channel to manage the control channel traffic. A reserved virtual LAN is used for the control channel traffic. Multiple Shared Ethernet Adapters are configured in a high availability mode without a dedicated control channel adapter, and are supported in this configuration.

Shared Ethernet Adapter failover with load sharing can also be configured without using a dedicated control channel adapter.

Scenario: Configuring Network Interface Backup in AIX client logical partitions without VLAN tagging

Use this scenario to become familiar with using a Network Interface Backup (NIB) configuration in Virtual I/O clients that are running AIX logical partitions and are not configured for VLAN tagging.

About this task

Situation

In this scenario, you want to configure a highly available virtual environment for your bridged network using the NIB approach to access external networks from your Virtual I/O clients. You do not plan to use VLAN tagging in your network setup. This approach requires you to configure a second Ethernet adapter on a different VLAN for each client and requires a Link Aggregation adapter with NIB features. This configuration is available for AIX logical partitions.

Note: You can also configure Ethernet bonding on Linux logical partitions. For more information, see the documentation for the Linux operating system.

Typically, a Shared Ethernet Adapter failover configuration is the recommended configuration for most environments because it supports environments with or without VLAN tagging. Also, the NIB configuration is more complex than a Shared Ethernet Adapter failover configuration because it must be implemented on each of the clients.

Note: In an NIB configuration, support for VLAN tagging is available only in the case where adapters that are configured under the NIB configuration, are configured under separate virtual switches. For example, when there are multiple virtual switches in an NIB configuration, you can add the primary adapter that is configured on *vswitch1* on *VLAN 20*, and the backup adapter that is configured on *vswitch2* also on *VLAN 20*.

However, Shared Ethernet Adapter failover was not available prior to Version 1.2 of Virtual I/O Server, and NIB was the only approach to a highly available virtual environment. Also, you might consider that in an NIB configuration you can distribute clients over both Shared Ethernet Adapters in such a way that half of them will use the first Shared Ethernet Adapter and the other half will use the second Shared Ethernet Adapter as primary adapter.

Objective

Create a virtual Ethernet environment using a Network Interface Backup configuration as depicted in the following figure.

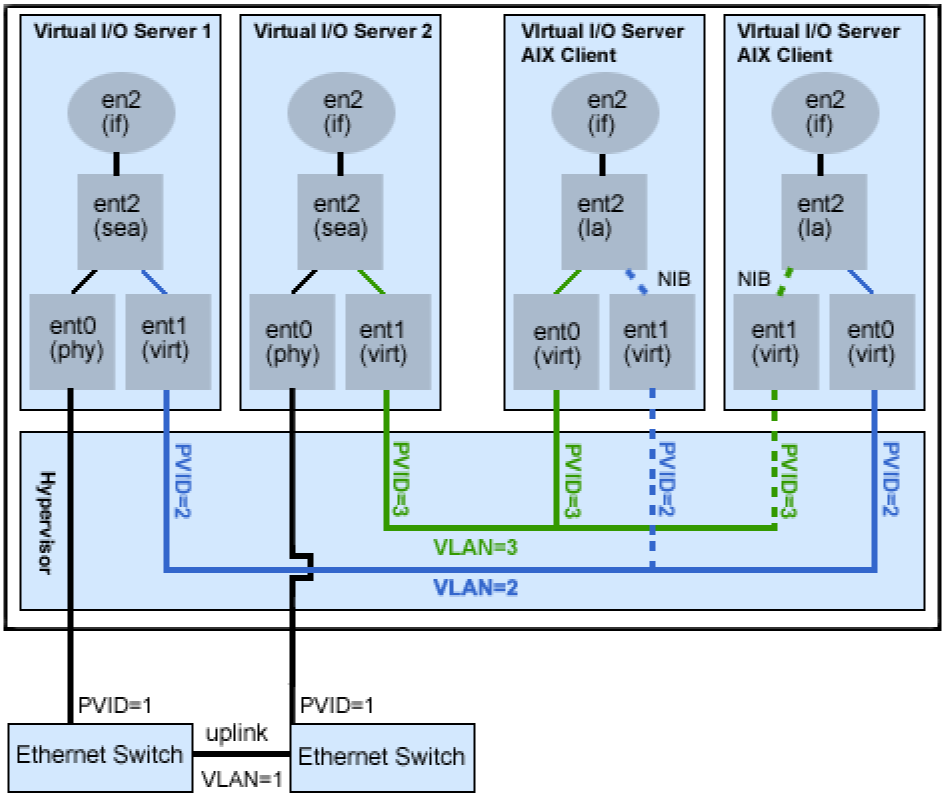

Prerequisites and assumptions

Before completing the configuration tasks, review the following prerequisites and assumptions.

- The Hardware Management Console (HMC) is already set up. For more information about Installing and configuring the HMC see Installing and configuring the Hardware Management Console.
- Two separate Virtual I/O Server logical partitions have been created and the Virtual I/O Server has been installed in each logical partition. See the instructions in "Installing the Virtual I/O Server and client logical partitions" on page 90.
- You have created the remaining logical partitions that you want added to the network configuration.
- Each Virtual I/O Server logical partition has an available physical Ethernet adapter assigned to it.

- You have IP addresses for all logical partitions and systems that will be added to the configuration.

Configuration tasks

Using the figure as a guide, complete the following tasks to configure the NIB virtual environment.

Procedure

1. Create a LAN connection between the Virtual I/O Servers and the external network:
   a) Configure a Shared Ethernet Adapter on the primary Virtual I/O Server that bridges traffic between the virtual Ethernet and the external network. See "Configuring a Shared Ethernet Adapter with the Virtual I/O Server command-line interface" on page 179.
   b) Configure a Shared Ethernet Adapter on the second Virtual I/O Server, as in step 1.
2. For each client logical partition, use the HMC to create a virtual Ethernet whose PVID matches the PVID of the primary Virtual I/O Server. This will be used as the primary adapter.
3. For each client logical partition, use the HMC to create a second virtual Ethernet whose PVID matches the PVID of the second (backup) Virtual I/O Server. This will be used as the backup adapter.
4. Create the Network Interface Backup setup using a Link Aggregation configuration. To create this configuration, follow the procedure Configuring an Etherchannel in the IBM Power Systems and AIX Information Center. Make sure that you specify the following items:
   a) Select the primary Ethernet Adapter.
   b) Select the Backup Adapter.
   c) Specify the Internet Address to Ping. Select the IP address or host name of a host outside of the Virtual I/O Server system that NIB will continuously ping to detect Virtual I/O Server failure.

Results

Note: Keep in mind, when you configure NIB with two virtual Ethernet adapters, the internal networks used must stay separated in the hypervisor. You must use different PVIDs for the two adapters in the client and cannot use additional VIDs on them.

Scenario: Configuring Multi-Path I/O for AIX client logical partitions

Multi-Path I/O (MPIO) helps provide increased availability of virtual Small Computer Serial Interface (SCSI) resources by providing redundant paths to the resource. This topic describes how to set up Multi-Path I/O for AIX client logical partitions.

Before you begin

To provide MPIO to AIX client logical partitions, you must have two Virtual I/O Server logical partitions configured on your system. This procedure assumes that the disks are already allocated to both the Virtual I/O Server logical partitions involved in this configuration.

Note: You can also configure MPIO on Linux logical partitions. For more information, see the documentation for the Linux operating system.

About this task

To configure MPIO, follow these steps. In this scenario, hdisk5 in the first Virtual I/O Server logical partition, and hdisk7 in the second Virtual I/O Server logical partition, are used in the configuration.

The following figure shows the configuration that is to be completed during this scenario.

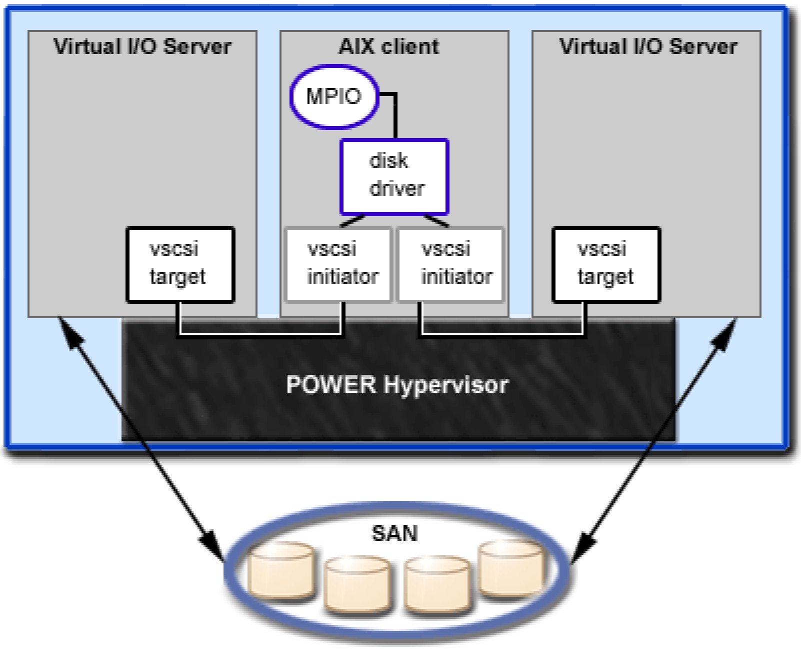

Using the preceding figure as a guide, follow these steps:

Procedure

1. Using the HMC, create SCSI server adapters on the two Virtual I/O Server logical partitions.
2. Using the HMC, create two virtual client SCSI adapters on the client logical partitions, each mapping to one of the Virtual I/O Server logical partitions.
3. On either of the Virtual I/O Server logical partitions, determine which disks are available by typing `lsdev -type disk`. Your results look similar to the following:

   ```
   name            status     description

   hdisk3          Available  MPIO Other FC SCSI Disk Drive
   hdisk4          Available  MPIO Other FC SCSI Disk Drive
   hdisk5          Available  MPIO Other FC SCSI Disk Drive
   ```

   Select which disk that you want to use in the MPIO configuration. In this scenario, hdisk5 is selected.
4. Determine the ID of the disk that you have selected. For instructions, see "Identifying exportable disks" on page 120. In this scenario, the disk does not have an IEEE volume attribute identifier or a unique identifier (UDID). Hence, determine the physical identifier (PVID) by running the `lspv hdisk5` command. Your results look similar to the following:

   ```
   hdisk5          00c3e35ca560f919                    None
   ```

   The second value is the PVID. In this scenario, the PVID is 00c3e35ca560f919. Note this value.
5. List the attributes of the disk on the first Virtual I/O Server by using the **lsdev** command. In this scenario, type `lsdev -dev hdisk5 -attr`. Your results look similar to the following:

   ```
   ..
   lun_id          0x5463000000000000              Logical Unit Number ID          False
   ..
   ..
   pvid            00c3e35ca560f9190000000000000000 Physical volume identifier      False
   ..
   reserve_policy  single_path                     Reserve Policy                  True
   ```

   Note the values for lun_id and reserve_policy. If the reserve_policy attribute is set to anything other than `no_reserve`, then you must change it. Set the reserve_policy to `no_reserve` by typing `chdev -dev hdisk`*x* `-attr reserve_policy=no_reserve`.

6. On the second Virtual I/O Server logical partition, list the physical volumes by typing `lspv`. In the output, locate the disk that has the same PVID as the disk identified previously. In this scenario, the PVID for hdisk7 matched:

   ```
   hdisk7          00c3e35ca560f919                    None
   ```

   Tip: Although the PVID values must be identical, the disk numbers on the two Virtual I/O Server logical partitions might vary.

7. Determine if the reserve_policy attribute is set to no_reserve using the **lsdev** command. In this scenario, type `lsdev -dev hdisk7 -attr`. You see results similar to the following:

   ```
   ..
   lun_id          0x5463000000000000              Logical Unit Number ID          False
   ..
   pvid            00c3e35ca560f9190000000000000000 Physical volume identifier      False
   ..
   reserve_policy  single_path                     Reserve Policy
   ```

   If the reserve_policy attribute is set to anything other than `no_reserve`, you must change it. Set the reserve_policy to `no_reserve` by typing `chdev -dev hdisk`*x* `-attr reserve_policy=no_reserve`.

8. On both Virtual I/O Server logical partitions, use the **mkvdev** to create the virtual devices. In each case, use the appropriate hdisk value. In this scenario, type the following commands:

   - On the first Virtual I/O Server logical partition, type `mkvdev -vdev hdisk5 -vadapter vhost5 -dev vhdisk5`
   - On the second Virtual I/O Server logical partition, type `mkvdev -vdev hdisk7 -vadapter vhost7 -dev vhdisk7`

   The same LUN is now exported to the client logical partition from both Virtual I/O Server logical partitions.

9. AIX can now be installed on the client logical partition. For instructions on installing AIX, see Installing AIX in a Partitioned Environment in the IBM Power Systems and AIX Information Center.

10. After you have installed AIX on the client logical partition, check for MPIO by running the following command:

    ```
    lspath
    ```

    You see results similar to the following:

    ```
    Enabled hdisk0 vscsi0
    Enabled hdisk0 vscsi1
    ```

    If one of the Virtual I/O Server logical partitions fails, the results of the **`lspath`** command look similar to the following:

    ```
    Failed  hdisk0 vscsi0
    Enabled hdisk0 vscsi1
    ```

    Unless a health check is enabled, the state continues to show `Failed` even after the disk has recovered. To have the state updated automatically, type `chdev -l hdisk`*x* `-a hcheck_interval=60 -P`. The client logical partition must be rebooted for this change to take effect.

Planning for the Virtual I/O Server

Use this topic to help gain an understanding of what to consider when you plan for the Virtual I/O Server.

Specifications required to create the Virtual I/O Server

This topic defines the range of configuration possibilities, including the minimum number of resources that are needed and the maximum number of resources that are allowed to create the Virtual I/O Server (VIOS).

To activate the VIOS, the PowerVM Editions hardware feature is required. A logical partition with enough resources to share with other logical partitions is required. The following is a list of minimum hardware requirements that must be available to create the VIOS.

Table 22. Resources that are required

Resource	**Requirement**
Hardware Management Console	The HMC is required to create the logical partition and assign resources.
Storage adapter	The server logical partition needs at least one storage adapter.
Physical disk	The disk must be at least 30 GB. This disk can be shared.
Ethernet adapter	If you want to route network traffic from virtual Ethernet adapters to a Shared Ethernet Adapter, you need an Ethernet adapter.
Memory	For POWER7, POWER8, or POWER9 processor-based systems, at least 768 MB of memory is required.
Processor	At least 0.05 processor use is required.

The following table defines the limitations for storage management.

Table 23. Limitations for storage management

Category	**Limit**
Volume groups	4096 per system
Physical volumes	1024 per volume group
Physical partitions	1024 per volume group
Logical volumes	1024 per volume group
Logical partitions	No limit

Limitations and restrictions of the Virtual I/O Server configuration

Learn about Virtual I/O Server (VIOS) configuration limitations.

Consider the following when you implement virtual Small Computer Serial Interface (SCSI):

- Virtual SCSI supports the following connection standards for backing devices: Fibre Channel, SCSI, SCSI RAID, iSCSI, SAS, SATA, USB, and IDE.
- The SCSI protocol defines mandatory and optional commands. While virtual SCSI supports all of the mandatory commands, not all of the optional commands are supported.
- There might be utilization implications when you use virtual SCSI devices. Because the client/server model is made up of layers of function, using virtual SCSI can consume additional processor cycles when processing I/O requests.
- The VIOS is a dedicated logical partition to be used only for VIOS operations. Other applications cannot run in the VIOS logical partition.

- If there is a resource shortage, performance degradation might occur. If a VIOS is serving many resources to other logical partitions, ensure that enough processor power is available. In case of high workload across virtual Ethernet adapters and virtual disks, logical partitions might experience delays in accessing resources.
- Logical volumes and files exported as virtual SCSI disks are always configured as single path devices on the client logical partition.
- Logical volumes or files exported as virtual SCSI disks that are part of the root volume group (rootvg) are not persistent if you reinstall the VIOS. However, they are persistent if you update the VIOS to a new service pack. Therefore, before reinstalling the VIOS, ensure that you back up the corresponding clients' virtual disks. When exporting logical volumes, it is best to export logical volumes from a volume group other than the root volume group. When exporting files, it is best to create file storage pools and the virtual media repository in a parent storage pool other than the root volume group.

Consider the following when you implement virtual adapters:

- Only Ethernet adapters can be shared. Other types of network adapters cannot be shared.
- IP forwarding is not supported on the VIOS.
- The maximum number of virtual adapters can be any value in the range 2 - 65,536. However, if you set the maximum number of virtual adapters to a value higher than 1024, the logical partition might fail to activate or the server firmware might require more system memory to manage the virtual adapters.

Consider the following when you increase the virtual I/O slot limit:

- The maximum number of virtual I/O slots supported on AIX, IBM i, and Linux partition is up to 32767.
- The maximum number of virtual adapters can be any value in the range 2 - 32767. However, higher maximum values require more system memory to manage the virtual adapters.

For more information about operating systems that run on client logical partitions and that are supported by the Virtual I/O Server (VIOS), see System software maps.

Capacity planning

This topic includes capacity-planning considerations for the Virtual I/O Server, including information about hardware resources and limitations.

Client logical partitions might use virtual devices, dedicated devices, or a combination of both. Before you begin to configure and install the Virtual I/O Server and client logical partitions, plan what resources each logical partition uses. Throughput requirements and overall workload must be considered when you decide whether to use virtual or dedicated devices and when you allocate resources to the Virtual I/O Server. Compared to dedicated Small Computer Serial Interface (SCSI) disks, virtual SCSI disks might achieve similar throughput numbers depending on several factors, including workload and virtual SCSI resources. However, virtual SCSI devices generally have higher processor utilization when compared with directly attached storage.

Planning for virtual SCSI

Find capacity-planning and performance information for virtual Small Computer Serial Interface (SCSI).

Different I/O subsystems have different performance qualities, as does virtual SCSI. This section discusses the performance differences between physical and virtual I/O. The following topics are described in this section:

Virtual SCSI latency

Find information about virtual Small Computer Serial Interface (SCSI) latency.

I/O latency is the amount of time that passes between the initiation and completion of a disk I/O operation. For example, consider a program that performs 1000 random disk I/O operations, one at a time. If the time to complete an average operation is 6 milliseconds, the program runs in no fewer than 6 seconds. However, if the average response time is reduced to 3 milliseconds, the run time might be reduced by 3 seconds. Applications that are multithreaded or use asynchronous I/O might be less sensitive to latency, but in most circumstances, lesser latency can help improve performance.

Because virtual SCSI is implemented as a client and server model, there is some latency that does not exist with directly attached storage. The latency might range from 0.03 to 0.06 milliseconds per I/O operation depending primarily on the block size of the request. The average latency is comparable for both physical disk and logical volume-backed virtual drives. The latency experienced when you use a Virtual I/O Server in a shared-processor logical partition can be higher and more variable than using a Virtual I/O Server in a dedicated logical partition. For more information about the performance differences between dedicated logical partitions and shared-processor logical partitions, see "Virtual SCSI sizing considerations" on page 70.

The following table identifies latency (in milliseconds) for different block-size transmissions on both physical disk and logical-volume-backed virtual SCSI disks.

Table 24. Increase in disk I/O response time based on block size (in milliseconds)

Backing type	**4 K**	**8 K**	**32 K**	**64 K**	**128 K**
Physical disk	0.032	0.033	0.033	0.040	0.061
Logical volume	0.035	0.036	0.034	0.040	0.063

The average disk-response time increases as the block size increases. The latency increases for a virtual SCSI operation are relatively greater on smaller block sizes because of their shorter response time.

Virtual SCSI bandwidth

View information about virtual Small Computer Serial Interface (SCSI) bandwidth.

I/O bandwidth is the maximum amount of data that can be read or written to a storage device in a unit of time. Bandwidth can be measured from a single thread or from a set of threads that run concurrently. Although many customer applications are more sensitive to latency than bandwidth, bandwidth is crucial for many typical operations, such as backing up and restoring persistent data.

The following table compares the results of bandwidth tests for virtual SCSI and physical I/O performance. In the tests, a single thread operates sequentially on a constant file that is 256 MB with a Virtual I/O Server running in a dedicated partition. More I/O operations are issued when reading or writing to the file by using a small block size as compared to a larger block size. The test was conducted by using a storage server with feature code 6239 (type 5704/0625) and a 2-gigabit Fibre Channel adapter attached to one RAID0 LUN that is composed of five physical disks from a DS4400 disk system (formerly a FAStT700). The table shows the comparison of measured bandwidth in megabytes per second (MB/s) by using virtual SCSI and local attachment for reads with varying block sizes of operations. The difference between virtual I/O and physical I/O in these tests is attributable to the increased latency when using virtual I/O. Because of the larger number of operations, the bandwidth measured with small block sizes is lesser than with large block sizes.

Table 25. Physical and virtual SCSI bandwidth comparison (in MB/s)

I/O type	**4 K**	**8 K**	**32 K**	**64 K**	**128 K**
Virtual	20.3	35.4	82.6	106.8	124.5
Physical	24.3	41.7	90.6	114.6	132.6

Virtual SCSI sizing considerations

Understand the processor and memory-sizing considerations when you implement virtual Small Computer Serial Interface (SCSI).

When you are designing and implementing a virtual SCSI application environment, consider the following sizing issues:

- The amount of memory that is allocated to the Virtual I/O Server
- The processor entitlement of the Virtual I/O Server

- Whether the Virtual I/O Server is run as a shared-processor logical partition or as a dedicated processor logical partition
- The maximum transfer size limitation for physical devices and AIX clients and AIX clients

The processor impacts of using virtual I/O on the client are insignificant. The processor cycles run on the client to perform a virtual SCSI I/O operation are comparable to that of a locally attached I/O device. Thus, there is no increase or decrease in sizing on the client logical partition for a known task. These sizing techniques do not anticipate combining the function of shared Ethernet with the virtual SCSI server. If the two are combined, consider adding resources to account for the shared Ethernet activity with virtual SCSI.

Virtual SCSI sizing using dedicated processor logical partitions

The amount of processor entitlement required for a virtual SCSI server is based on the maximum I/O rates required of it. Because virtual SCSI servers do not normally run at maximum I/O rates all of the time, the use of surplus processor time is potentially wasted when you use dedicated processor logical partitions. In the first of the following sizing methodologies, you need a good understanding of the I/O rates and I/O sizes required of the virtual SCSI server. In the second, size the virtual SCSI server based on the I/O configuration.

The sizing methodology used is based on the observation that the processor time required to perform an I/O operating on the virtual SCSI server is fairly constant for a given I/O size. It is a simplification to make this statement, because different device drivers have subtly varying efficiencies. However, under most circumstances, the I/O devices supported by the virtual SCSI server are sufficiently similar. The following table shows approximate cycles per second for both physical disk and logical volume operations on a 1.65 Ghz processor. These numbers are measured at the physical processor; simultaneous multithreading (SMT) operation is assumed. For other frequencies, scaling by the ratio of the frequencies (for example, 1.5 Ghz = 1.65 Ghz / 1.5 Ghz × cycles per operation) is sufficiently accurate to produce a reasonable sizing.

Table 26. Approximate cycles per second on a 1.65 Ghz logical partition

Disk type	**4 KB**	**8 KB**	**32 KB**	**64 KB**	**128 KB**
Physical disk	45,000	47,000	58,000	81,000	120,000
Logical volume	49,000	51,000	59,000	74,000	105,000

Consider a Virtual I/O Server that uses three client logical partitions on physical disk-backed storage. The first client logical partition requires a maximum of 7,000 8-KB operations per second. The second client logical partition requires a maximum of 10,000 8-KB operations per second. The third client logical partition requires a maximum of 5,000 128-KB operations per second. The number of 1.65 Ghz processors for this requirement is approximately ((7,000 × 47,000 + 10,000 × 47,000 + 5,000 × 120,000) / 1,650,000,000) = 0.85 processors, which rounds up to a single processor when you use a dedicated processor logical partition.

If the I/O rates of the client logical partitions are not known, you can size the Virtual I/O Server to the maximum I/O rate of the storage subsystem attached. The sizing might be biased toward small I/O operations or large I/O operations. Sizing to maximum capacity for large I/O operations balance the processor capacity of the Virtual I/O Server to the potential I/O bandwidth of the attached I/O. The negative aspect of this sizing methodology is that, in nearly every case, more processor entitlement is assigned to the Virtual I/O Server than it typically consumes.

Consider a case in which a Virtual I/O Server manages 32 physical SCSI disks. A maximum limit of processors required can be established based on assumptions about the I/O rates that the disks can achieve. If it is known that the workload is dominated by 8096-byte operations that are random, then assume that each disk is capable of approximately 200 disk I/O operations per second (15k rpm drives). At peak, the Virtual I/O Server would need to serve approximately 32 disks × 200 I/O operations per second × 47,000 cycles per operation, resulting in a requirement for approximately 0.19 processor performance. Viewed another way, a Virtual I/O Server running on a single processor must be capable of supporting more than 150 disks doing 8096 byte random I/O operations.

Alternatively, if the Virtual I/O Server is sized for maximum bandwidth, the calculation results in a higher processor requirement. The difference is that maximum bandwidth assumes sequential I/O. Because disks are more efficient when they are performing large, sequential I/O operations than they are when performing small, random I/O operations, a higher number of I/O operations per second can be performed. Assume that the disks are capable of 50 MB per second when doing 128 KB I/O operations. That situation implies each disk might average 390 disk I/O operations per second. Thus, the amount of processing power necessary to support 32 disks, each doing 390 I/O operations per second with an operation cost of 120,000 cycles (32 × 390 × 120,000 / 1,650,000,000) results in approximately 0.91 processors. Consequently, a Virtual I/O Server running on a single processor must be capable of driving approximately 32 fast disks to maximum throughput.

Virtual SCSI server sizing using shared processor logical partitions

Defining virtual SCSI servers in shared processor logical partitions allows more specific processor resource sizing and potential recovery of unused processor time by uncapped logical partitions. However, using shared-processor logical partitions for virtual SCSI servers can frequently increase I/O response time and make for more complex processor entitlement sizings.

The sizing methodology must be based on the same operation costs for dedicated logical partition I/O servers, with added entitlement for running in shared-processor logical partitions. Configure the Virtual I/O Server as uncapped, so that, if the Virtual I/O Server is undersized, there is opportunity to get more processor time to serve I/O operations.

Because I/O latency with virtual SCSI can vary due to a number of conditions, consider the following if a logical partition has high I/O requirements:

- Configure the logical partition with physical I/O if the configuration allows.
- In most cases, the Virtual I/O Server logical partition can use a shared, uncapped processor.

Virtual SCSI server memory sizing

Memory sizing in virtual SCSI is simplified because there is no caching of file data in the memory of the virtual SCSI server. Because there is no data caching, the memory requirements for the virtual SCSI server are fairly modest. With large I/O configurations and very high data rates, a 1 GB memory allocation for the virtual SCSI server is likely to be sufficient. For low I/O rate situations with a few attached disks, 512 MB most likely suffices.

Virtual SCSI maximum transfer size limitation

If you add another virtual target device to the virtual SCSI server adapter and the new virtual target device has a smaller maximum transfer size than the other configured devices on that adapter, the Virtual I/O Server does not show a new virtual device to the client. At the time the virtual target device is created, the Virtual I/O Server displays a message stating that the new target device will not be visible to the client until you reboot the client.

To display the maximum transfer size of a physical device, use the following command: `lsdev -attr max_transfer -dev hdiskN`

Planning for Shared Ethernet Adapters

Use this section to find capacity-planning and performance information for Shared Ethernet Adapter. This section contains planning information and performance considerations for using **Shared Ethernet Adapters** on the Virtual I/O Server.

Network requirements

This topic includes information that you need to accurately size your Shared Ethernet Adapter environment.

To plan for using **Shared Ethernet Adapters**, you must determine your network needs. This section gives overview information of what must be considered when sizing the Shared Ethernet Adapter environment. Sizing the Virtual I/O Server for the Shared Ethernet Adapter involves the following factors:

- Defining the target bandwidth (MB per second), or transaction rate requirements (operations per second). The target performance of the configuration must be determined from your workload requirements.
- Defining the type of workload (streaming or transaction oriented).
- Identifying the maximum transmission unit (MTU) size that is to be used (1500 or jumbo frames).
- Determining if the Shared Ethernet Adapter runs in a threaded or nonthreaded environment.
- Knowing the throughput rates that various Ethernet adapters can provide (see Adapter selection).
- Knowing the processor cycles required per byte of throughput or per transaction (see Processor allocation).

Bandwidth requirement

The primary consideration is determining the target bandwidth on the physical Ethernet adapter of the Virtual I/O Server. This determines the rate that data can be transferred between the Virtual I/O Server and the client logical partitions. After the target rate is known, the correct type and number of network adapters can be selected. For example, Ethernet adapters of various speeds might be used. One or more adapters might be used on individual networks, or they might be combined by using Link Aggregation (or Etherchannel).

Workload type

The type of workload to be performed must be considered, whether it is streaming of data for workloads such as file transfer, data backup, or small transaction workloads, such as remote procedure calls. The streaming workload consists of large, full-sized network packets and associated small, TCP acknowledgment packets. Transaction workloads typically involve smaller packets or might involve small requests, such as a URL, and a larger response, such as a web page. A Virtual I/O Server needs to frequently support streaming and small packet I/O during various periods of time. In that case, approach the sizing from both models.

MTU size

The MTU size of the network adapters must also be considered. The standard Ethernet MTU is 1500 bytes. Gigabit Ethernet and 10-gigabit Ethernet can support 9000-byte MTU jumbo frames. Jumbo frames might reduce the processor cycles for the streaming types of workloads. However, for small workloads, the larger MTU size might not help reduce processor cycles.

Threaded or nonthreaded environment

Use threaded mode when virtual Small Computer Serial Interface (SCSI) is to be run on the same Virtual I/O Server logical partition as Shared Ethernet Adapter. Threaded mode helps ensure that virtual SCSI and the Shared Ethernet Adapter can share the processor resource appropriately. However, threading increases instruction-path length, which uses additional processor cycles. If the Virtual I/O Server logical partition will be dedicated to running shared Ethernet devices (and associated virtual Ethernet devices) only, the adapters must be configured with threading disabled. For more information, see "Processor allocation" on page 76.

Adapter throughput

Knowing the throughput capability of different Ethernet adapters can help you determine which adapters to use as **Shared Ethernet Adapters** and how many adapters to use. For more information, see "Adapter selection" on page 74.

Processor entitlement

You must determine how much processor power is required to move data through the adapters at the required rate. Networking device drivers are typically processor-intensive. Small packets can come in at a faster rate and use more processor cycles than larger packet workloads. Larger packet workloads are typically limited by network wire bandwidth and come in at a slower rate, thus requiring less processor power than small packet workloads for the amount of data transferred.

Adapter selection

Use this section to find the attributes and performance characteristics of various types of Ethernet adapters to help you select which adapters to use in your environment.

This section provides approximate throughput rates for various Ethernet adapters set at various MTU sizes. Use this information to determine which adapters are needed to configure a Virtual I/O Server. To make this determination, you must know the required throughput rate of the client logical partitions.

Following are general guidelines for network throughput. These numbers are not specific, but they can serve as a general guideline for sizing. In the following tables, the 100 MB, 1 GB, and 10 GB speeds are rounded down for estimating.

Table 27. Simplex (one direction) streaming rates

Adapter speed	Approximate throughput rate
10 Mb Ethernet	1 MB/second
100 Mb Ethernet	10 MB/second
1000 Mb Ethernet (GB Ethernet)	100 MB/second
10000 Mb Ethernet (10 GB Ethernet, Host Ethernet Adapter or Integrated Virtual Ethernet)	1000 MB/second

Table 28. Full duplex (two direction) streaming rates on full duplex network

Adapter speed	Approximate throughput rate
10 Mb Ethernet	2 MB/second
100 Mb Ethernet	20 MB/second
1000 Mb Ethernet (Gb Ethernet)	150 MB/second
10000 Mb Ethernet (10 Gb Ethernet, Host Ethernet Adapter or Integrated Virtual Ethernet)	1500 MB/second

The following tables list maximum network payload speeds, which are user payload data rates that can be obtained by sockets-based programs for applications that are streaming data. The rates are a result of the network bit rate, MTU size, excessive increase in physical levels (such as interframe gaps and preamble bits), data link headers, and TCP/IP headers. A gigahertz-speed processor is assumed. These numbers are optimal for a single LAN. If your network traffic is going through additional network devices, your results might vary.

In the following tables, raw bit rate is the physical media bit rate and does not reflect interframe gaps, preamble bits, data link headers, and trailers. Interframe gaps, preamble bits, data link headers, and trailers can all reduce the effective usable bit rate of the wire.

Single direction (simplex) TCP streaming rates are rates that can be achieved by sending data from one machine to another in a memory-to-memory test. Full-duplex media can usually perform slightly better than half-duplex media because the TCP acknowledgment packets can flow without contending for the same wire that the data packets are flowing on.

Table 29. Single direction (simplex) TCP streaming rates

Network type	Raw bit rate (Mb)	Payload rate (Mb)	Payload rate (MB)
10 Mb Ethernet, Half Duplex	10	6	0.7
10 Mb Ethernet, Full Duplex	10 (20 Mb full duplex)	9.48	1.13
100 Mb Ethernet, Half Duplex	100	62	7.3
100 Mb Ethernet, Full Duplex	100 (200 Mb full duplex)	94.8	11.3
1000 Mb Ethernet, Full Duplex, MTU 1500	1000 (2000 Mb full duplex)	948	113
1000 Mb Ethernet, Full Duplex, MTU 9000	1000 (2000 Mb full duplex)	989	117.9
10000 Mb Ethernet, Full Duplex, Host Ethernet Adapter (or Integrated Virtual Ethernet) MTU 1500	10000	9479	1130
10000 Mb Ethernet, Full Duplex, Host Ethernet Adapter (or Integrated Virtual Ethernet) MTU 9000	10000	9899	1180

Full-duplex TCP streaming workloads have data streaming in both directions. Workloads that can send and receive packets concurrently can take advantage of full duplex media. Some media, for example Ethernet in half-duplex mode, cannot send and receive concurrently, thus they do not perform any better, and can usually degrade performance, when running duplex workloads. Duplex workloads do not increase at a full doubling of the rate of a simplex workload because the TCP acknowledgment packets returning from the receiver must now compete with data packets flowing in the same direction.

Table 30. Two direction (duplex) TCP streaming rates

Network type	Raw bit rate (Mb)	Payload rate (Mb)	Payload rate (MB)
10 Mb Ethernet, Half Duplex	10	5.8	0.7
10 Mb Ethernet, Full Duplex	10 (20 Mb full duplex)	18	2.2
100 Mb Ethernet, Half Duplex	100	58	7

Table 30. Two direction (duplex) TCP streaming rates (continued)

Network type	Raw bit rate (Mb)	Payload rate (Mb)	Payload rate (MB)
100 Mb Ethernet, Full Duplex	100 (200 Mb full duplex)	177	21.1
1000 Mb Ethernet, Full Duplex, MTU 1500	1000 (2000 Mb full duplex)	1470 (1660 peak)	175 (198 peak)
1000 Mb Ethernet, Full Duplex, MTU 9000	1000 (2000 Mb full duplex)	1680 (1938 peak)	200 (231 peak)
10000 Mb Ethernet, Host Ethernet Adapter (or Integrated Virtual Ethernet) Full Duplex, MTU 1500	10000	14680 (15099 peak)	1750 (1800 peak)
10000 Mb Ethernet, Host Ethernet Adapter (or Integrated Virtual Ethernet) Full Duplex, MTU 9000	10000	16777 (19293 pack)	2000 (2300 peak)

Notes:

1. Peak numbers represent optimal throughput with multiple TCP sessions running in each direction. Other rates are for a single TCP session.
2. 1000 MB Ethernet (gigabit Ethernet) duplex rates are for the PCI-X adapter in PCI-X slots.
3. Data rates are for TCP/IP by using the IPv4 protocol. Adapters with MTU set to 9000 have RFC 1323 enabled.

Processor allocation

This section contains processor-allocation guidelines for both dedicated processor logical partitions and shared processor logical partitions.

Because Ethernet running MTU size of 1500 bytes consumes more processor cycles than Ethernet running Jumbo frames (MTU 9000), the guidelines are different for each situation. In general, the processor utilization for large packet workloads on jumbo frames is approximately half that required for MTU 1500.

If MTU is set to 1500, provide one processor (1.65 Ghz) per Gigabit Ethernet adapter to help reach maximum bandwidth. This equals ten 100-Mb Ethernet adapters if you are using smaller networks. For smaller transaction workloads, plan to use one full processor to drive the Gigabit Ethernet workload to maximum throughput. For example, if 2 Gigabit Ethernet adapters are to be used, allocate up to two processors to the logical partition.

If MTU is set to 9000 (jumbo frames), provide 50% of one processor (1.65 Ghz) per Gigabit Ethernet adapter to reach maximum bandwidth. Small packet workloads must plan to use one full processor to drive the Gigabit Ethernet workload. Jumbo frames have no effect on the small packet workload case.

Shared Ethernet Adapter using a dedicated processor logical partition

The sizing provided is divided into two workload types: TCP streaming and TCP request and response. Both MTU 1500 and MTU 9000 networks were used in the sizing, which is provided in terms of machine cycles per byte of throughput for streaming or per transaction for request/response workloads.

The data in the following tables was derived by using the following formula:

(number of processors × processor_utilization × processor clock frequency) / Throughput rate in bytes per second or transaction per second = cycles per Byte or transaction.

For the purposes of this test, the numbers were measured on a logical partition with one 1.65 Ghz processor with simultaneous multi-threading (SMT) enabled.

For other processor frequencies, the numbers in these tables can be scaled by the ratio of the processor frequencies for approximate values to be used for sizing. For example, for a 1.5 Ghz processor speed, use 1.65/1.5 × cycles per byte value from the table. This example would result in a value of 1.1 times the value in the table, thus requiring 10% more cycles to adjust for the 10% slower clock rate of the 1.5 Ghz processor.

To use these values, multiply your required throughput rate (in bytes or transactions) by the cycles per byte value in the following tables. This result gives you the required machine cycles for the workload for a 1.65 Ghz speed. Then adjust this value by the ratio of the actual machine speed to this 1.65 Ghz speed. To find the number of processors, divide the result by 1,650,000,000 cycles (or the cycles rate if you adjusted to a different speed machine). You would need the resulting number of processors to drive the workload.

For example, if the Virtual I/O Server must deliver 200 MB of streaming throughput, the following formula would be used:

200 × 1024 × 1024 × 11.2 = 2,348,810,240 cycles / 1,650,000,000 cycles per processor = 1.42 processors.

In round numbers, it would require 1.5 processors in the Virtual I/O Server to handle this workload. Such a workload can then be handled by either a logical partition using two dedicated processors or by a logical partition using 1.5-processor shared processors.

The following tables show the machine cycles per byte for a TCP-streaming workload.

Table 31. Shared Ethernet with threading option enabled

Type of Streaming	MTU 1500 rate and processor utilization	MTU 1500, cycles per byte	MTU 9000 rate and processor utilization	MTU 9000, cycles per byte
Simplex	112.8 MB at 80.6% processor	11.2	117.8 MB at 37.7% processor	5
Duplex	162.2 MB at 88.8% processor	8.6	217 MB at 52.5% processor	3.8

Table 32. Shared Ethernet with threading option disabled

Type of Streaming	MTU 1500 rate and processor utilization	MTU 1500, cycles per byte	MTU 9000 rate and processor utilization	MTU 9000, cycles per byte
Simplex	112.8 MB at 66.4% processor	9.3	117.8 MB at 26.7% processor	3.6
Duplex	161.6 MB at 76.4% processor	7.4	216.8 MB at 39.6% processor	2.9

The following tables show the machine cycles per transaction for a request and response workload. A transaction is defined as a round-trip request and reply size.

Table 33. Shared Ethernet with threading option enabled

Size of transaction	Transactions per second and Virtual I/O Server utilization	MTU 1500 or 9000, cycles per transaction
Small packets (64 bytes)	59,722 TPS at 83.4% processor	23,022
Large packets (1024 bytes)	51,956 TPS at 80% processor	25,406

Table 34. Shared Ethernet with threading option disabled

Size of transaction	Transactions per second and Virtual I/O Server utilization	MTU 1500 or 9000, cycles per transaction
Small packets (64 bytes)	60,249 TPS at 65.6% processor	17,956
Large packets (1024 bytes)	53,104 TPS at 65% processor	20,196

The preceding tables demonstrate that the threading option of the shared Ethernet adds approximately 16% – 20% more machine cycles per transaction for MTU 1500 streaming, and approximately 31% – 38% more machine cycles per transaction for MTU 9000. The threading option adds more machine cycles per transaction at lesser workloads due to the threads being started for each packet. At higher workload rates, like full duplex or the request and response workloads, the threads can run longer without waiting and being redispatched. You can configure the thread option for each shared Ethernet adapter by using the Virtual I/O Server commands. Disable the thread option if the shared Ethernet is running in a Virtual I/O Server logical partition by itself (without virtual Small Computer Serial Interface (SCSI) in the same logical partition).

You can enable or disable threading using the **-attr thread** option of the **`mkvdev`** command. To enable threading, use the `-attr thread=1` option. To disable threading, use the `-attr thread=0` option. For example, the following command disables threading for Shared Ethernet Adapter `ent1`:

```
mkvdev -sea ent1 -vadapter ent5 -default ent5 -defaultid 1 -attr thread=0
```

Sizing a Virtual I/O Server for shared Ethernet on a shared processor logical partition

Creating a shared-processor logical partition for a Virtual I/O Server can be done if the Virtual I/O Server is running slower-speed networks (for example 10/100 Mb) and a full processor logical partition is not needed. It is suggested that this be done only if the Virtual I/O Server workload is less than half a processor or if the workload is inconsistent. Configuring the Virtual I/O Server logical partition as uncapped might also allow it to use more processor cycles as needed to handle inconsistent throughput. For example, if the network is used only when other processors are idle, the Virtual I/O Server logical partition might be able to use other machine cycles and might be created with minimal processor to handle light workload during the day but the uncapped processor might use more machine cycles at night.

If you are creating a Virtual I/O Server in a shared-processor logical partition, add additional entitled processors as a sizing contingency.

Memory allocation

Find information about memory allocation and sizing.

In general, 512 MB of memory per logical partition is sufficient for most configurations. Enough memory must be allocated for the Virtual I/O Server data structures. Ethernet adapters and virtual devices use dedicated receive buffers. These buffers are used to store the incoming packets, which are then sent over the outgoing device.

A physical Ethernet adapter typically uses 4 MB for MTU 1500 or 16 MB for MTU 9000 for dedicated receive buffers for gigabit Ethernet. Other Ethernet adapters are similar. Virtual Ethernet, typically uses 6 MB for dedicated receive buffers. However, this number can vary based on workload. Each instance of a physical or virtual Ethernet would need memory for this number of buffers. In addition, the system has an mbuf buffer pool per processor that is used if additional buffers are needed. These mbufs typically occupy 40 MB.

Configuration requirements for shared memory

Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

System requirements

- The server must be a POWER7 processor-based server, or later.
- The server firmware must be at release 3.4.2, or later.
- The Hardware Management Console (HMC) must be at version 7 release 3.4.2, or later.
- The PowerVM Active Memory Sharing technology must be activated. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

Paging VIOS partition requirements

- VIOS partitions that provide access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool (hereafter referred to as *paging VIOS partitions*) cannot use shared memory. Paging VIOS partitions must use dedicated memory.
- Paging VIOS partitions must be at version 2.1.1, or later.
- On HMC-managed systems, consider configuring separate VIOS partitions as server partitions and paging VIOS partitions. For example, configure one VIOS partition to provide virtual resources to the shared memory partitions. Then, configure another VIOS partition as a paging VIOS partition.
- On HMC-managed systems, you can configure multiple VIOS partitions to provide access to paging space devices. However, you can assign only up to two of those VIOS partitions to the shared memory pool at any time.

Requirements for shared memory partitions

- Shared memory partitions must use shared processors.
- You can assign only virtual adapters to shared memory partitions. This means that you can dynamically add only virtual adapters to shared memory partitions. More specifically, the following table lists the virtual adapters that you can assign shared memory partitions.

Table 35. Virtual adapters that you can assign to shared memory partitions

AIX and Linux shared memory partitions	IBM i shared memory partitions
– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial adapters	– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial server adapters

Table 36. Virtual adapters that you can assign to shared memory partitions

Linux shared memory partitions
– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial adapters

You cannot assign Host Ethernet Adapters (HEA) or host connection adapters (HCA) to shared memory partitions.

- Shared memory partitions cannot use the barrier synchronization register.
- Shared memory partitions cannot use huge pages.
- AIX must be at version 6.1 Technology Level 3, or later, to run in a shared memory partition.
- IBM i must be at 6.1 with PTF SI32798, or later, to run in a shared memory partition.
- Virtual OptiConnect must not be enabled on IBM i shared memory partitions.
- SUSE Linux Enterprise Server must be at version 11, or later, to run in a shared memory partition.
- Red Hat Enterprise Server Version 6, or later, to run in a shared memory partition.
- You cannot configure IBM i logical partitions that provide virtual resources to other logical partitions as shared memory partitions. Logical partitions that provide virtual resources to other logical partitions in a shared memory environment must be VIOS partitions.

Requirements for paging space devices

- The paging space devices for AIX or Linux shared memory partitions must be at least the size of the maximum logical memory of the shared memory partition.
- The paging space devices for IBM i shared memory partitions must be at least the size of the maximum logical memory of the shared memory partition plus 8 KB for every megabyte. For example, if the maximum logical memory of the shared memory partition is 16 GB, its paging space device must be at least 16.125 GB.
- Paging space devices can be assigned only to one shared memory pool at a time. You cannot assign the same paging space device to a shared memory pool on one system and to another shared memory pool on another system at the same time.
- Paging space devices that are accessed by a single paging VIOS partition must meet the following requirements:
  - They can be physical or logical volumes.
  - They can be located in physical storage on the server or on a storage area network (SAN).
- Paging space devices that are accessed redundantly by two paging VIOS partitions must meet the following requirements:
  - They must be physical volumes.
  - They must be located on a SAN.
  - They must be configured with global IDs.
  - They must be accessible to both paging VIOS partitions.
  - The reserve attribute must be set to no reserve. (The VIOS automatically sets the reserve attribute to no reserve when you add the paging space device to the shared memory pool.)
- Physical volumes that are configured as paging space devices cannot belong to a volume group, such as the `rootvg` volume group.
- Logical volumes that are configured as paging space devices must be located in a volume group that is dedicated for paging space devices.
- Paging space devices must be available. You cannot use the physical volume or logical volume as a paging space device if it is already configured as a paging space device or virtual disk for another logical partition.
- Paging space devices cannot be used to boot a logical partition.
- After you assign a paging space device to the shared memory pool, you must manage the device by using the **Create/Modify Shared Memory Pool** wizard on the HMC. Do not change or remove the device by using other management tools.

Redundancy considerations

Redundancy options are available at several levels in the virtual I/O environment. Multipathing, mirroring, and RAID redundancy options exist for the Virtual I/O Server and some client logical partitions. Ethernet Link Aggregation (also called Etherchannel) is also an option for the client logical partitions, and the Virtual I/O Server provides Shared Ethernet Adapter failover. There is also support for node failover (PowerHA® SystemMirror®) for nodes that use virtual I/O resources.

This section contains information about redundancy for both the client logical partitions and the Virtual I/O Server. While these configurations help protect from the failure of one of the physical components, such as a disk or network adapter, they might cause the client logical partition to lose access to its devices if the Virtual I/O Server fails. The Virtual I/O Server can be made redundant by running a second instance of it in another logical partition. When you run two instances of the Virtual I/O Server, you can use LVM mirroring, multipath I/O, network interface backup, or multipath routing with dead gateway detection in the client logical partition to provide highly available access to virtual resources hosted in separate Virtual I/O Server logical partitions.

Client logical partitions

This topic includes redundancy considerations for client logical partitions. MPIO, PowerHA SystemMirror, and mirroring for the client logical partition are discussed.

Multipath I/O

View Multipath I/O (MPIO) information for client logical partitions.

Multiple virtual Small Computer Serial Interface (SCSI) or virtual Fibre Channel adapters in a client logical partition can access the same disk through multiple Virtual I/O Server logical partitions. This section describes a virtual SCSI multipath device configuration. If correctly configured, the client recognizes the disk as a multipath device. If you are using PowerVM Active Memory Sharing technology (or shared memory) or the Suspend/Resume feature, you can also use a multipath configuration to enable two paging VIOS logical partitions to access common paging space devices.

MPIO is not available for client logical partitions that run IBM i versions earlier that 6.1.1. Instead, you must use mirroring to create redundancy. For more information, see “Mirroring for client logical partitions” on page 82.

Not all virtual SCSI devices are capable of MPIO. To create an MPIO configuration, the exported device at the Virtual I/O Server must conform to the following rules:

- The device must be backed by a physical volume. Logical volume-backed virtual SCSI devices are not supported in an MPIO configuration.
- The device must be accessible from multiple Virtual I/O Server logical partitions.
- The device must be an MPIO-capable device.

  Note: MPIO-capable devices are those devices that contain a unique identifier (UDID) or IEEE volume identifier. For instructions about how to determine whether disks have a UDID or IEEE volume identifier, see “Identifying exportable disks” on page 120.

When you set up an MPIO configuration for virtual SCSI devices on the client logical partition, you must consider the reservation policy of the device on the Virtual I/O Server. To use an MPIO configuration at the client, none of the virtual SCSI devices on the Virtual I/O Server can be reserving the virtual SCSI device. Ensure the `reserve_policy` attribute of the device is set to `no_reserve`.

Failover is the only supported behavior for MPIO virtual SCSI disks on an AIX logical partition.

Related tasks

Setting the reserve policy attributes of a device
In some configurations, you must consider the reservation policy of the device on the Virtual I/O Server (VIOS).

Scenario: Configuring Multi-Path I/O for AIX client logical partitions

Multi-Path I/O (MPIO) helps provide increased availability of virtual Small Computer Serial Interface (SCSI) resources by providing redundant paths to the resource. This topic describes how to set up Multi-Path I/O for AIX client logical partitions.

Related reference

Configuration requirements for shared memory
Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

Mirroring for client logical partitions

Achieve mirroring for client logical partitions by using two virtual Small Computer Serial Interface (SCSI) adapters.

The client partition can mirror its logical volumes by using two virtual SCSI client adapters. Each of these adapters must be assigned to separate Virtual I/O Server partitions. The two physical disks are each attached to a separate Virtual I/O Server partition and made available to the client partition through a virtual SCSI server adapter. This configuration protects virtual disks in a client partition against the failure of any of the following:

- One physical disk
- One physical adapter
- One Virtual I/O Server

The performance of your system might be impacted when you use a RAID 1 configuration.

PowerHA SystemMirror in the Virtual I/O Server

Learn about PowerHA SystemMirror in the Virtual I/O Server.

PowerHA SystemMirror supports certain configurations that use the Virtual I/O Server, virtual Small Computer Serial Interface (SCSI), and virtual networking capabilities. For the most recent support and configuration information, see the IBM PowerHA SystemMirror for AIX website. For more information on PowerHA SystemMirror documentation, see PowerHA SystemMirror for AIX.

For IBM i client partitions, you must use mirroring to create redundancy. For more information, see "Mirroring for client logical partitions" on page 82.

PowerHA SystemMirror and virtual SCSI

Be aware of the following considerations when you implement PowerHA SystemMirror and virtual SCSI:

- The volume group must be defined as Enhanced Concurrent Mode. Enhanced Concurrent Mode is the preferred mode for sharing volume groups in PowerHA SystemMirror clusters because volumes are accessible by multiple PowerHA SystemMirror nodes. If file systems are used on the standby nodes, those file systems are not mounted until the point of failover. If shared volumes are accessed directly (without file systems) in Enhanced Concurrent Mode, these volumes are accessible from multiple nodes, and as a result, access must be controlled at a higher layer.
- If a cluster node accesses shared volumes by using virtual SCSI, all nodes in that cluster must also access the same shared volume. This means that disks cannot be shared between a logical partition by using virtual SCSI and a node directly accessing those disks.
- All volume group configuration and maintenance on these shared disks is done from the PowerHA SystemMirror nodes, not from the Virtual I/O Server.

PowerHA SystemMirror and virtual Ethernet

Be aware of the following considerations when you implement PowerHA SystemMirror and virtual Ethernet:

- IP Address Takeover (IPAT) by way of aliasing must be used. IPAT by way of Replacement and MAC Address Takeover are not supported.
- Avoid using the PowerHA SystemMirror PCI Hot Plug facility in a Virtual I/O Server environment. PCI Hot Plug operations are available through the Virtual I/O Server. When an PowerHA SystemMirror node

is using virtual I/O, the PowerHA SystemMirror PCI Hot Plug facility is not meaningful because the I/O adapters are virtual rather than physical.

- All virtual Ethernet interfaces defined to PowerHA SystemMirror must be treated as single-adapter networks. In particular, you must use the **ping_client_list** attribute to monitor and detect failure of the network interfaces.
- If the Virtual I/O Server has multiple physical interfaces on the same network, or if there are two or more PowerHA SystemMirror nodes that use the Virtual I/O Server in the same frame, PowerHA SystemMirror is not informed of, and does not react to, single physical interface failures. This does not limit the availability of the entire cluster because the Virtual I/O Server routes traffic around the failure.
- If the Virtual I/O Server has only a single physical interface on a network, failure of that physical interface is detected by PowerHA SystemMirror. However, that failure isolates the node from the network.

Link aggregation or Etherchannel devices

A link aggregation, or Etherchannel device, is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters that are aggregated can then act as a single Ethernet device. Link aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` adapters can be aggregated to the `ent3` adapter. The system considers these aggregated adapters as one adapter, and all adapters in the link aggregation device are given the same hardware address. Therefore, they are treated by remote systems as if they were one adapter.

Link aggregation can provide increased redundancy because individual links might fail. The link aggregation device can automatically fail over to another adapter in the device to maintain connectivity. For example, if the `ent0` adapter fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. The `ent0` adapter automatically returns to service on the link aggregation device when it recovers.

You can configure a Shared Ethernet Adapter to use a link aggregation, or Etherchannel, device as the physical adapter.

Shared Ethernet Adapter failover

Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server logical partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

A Shared Ethernet Adapter comprises a physical adapter (or several physical adapters grouped under a Link Aggregation device) and one or more virtual Ethernet adapters. It can provide layer 2 connectivity to multiple client logical partitions through the virtual Ethernet adapters.

The Shared Ethernet Adapter failover configuration involves two Shared Ethernet Adapters. One of the Shared Ethernet Adapters serves as the primary and one Shared Ethernet Adapter serves as the backup. There must be only one pair of Shared Ethernet Adapters per vSwitch/PVID. The VLAN IDs on this pair of Shared Ethernet Adapters must not be on any other Shared Ethernet Adapters on the vSwitch. The priority value that is given to the virtual Ethernet adapters during their creation is used to determine which Shared Ethernet Adapter serves as the primary and which Shared Ethernet Adapter serves as the backup. The Shared Ethernet Adapter that has the virtual Ethernet configured with the numerically lesser priority value will be used preferentially as the primary adapter. For the purpose of communicating between themselves to determine when a failover should take place, Shared Ethernet Adapters in failover mode use a VLAN dedicated for such traffic, called the *control channel*. For this reason, a virtual Ethernet (created with a PVID that is unique on the system) must be specified as the control channel virtual Ethernet when each Shared Ethernet Adapter is created in failover mode. Using the control channel, the backup Shared Ethernet Adapter is notified when the primary adapter fails, and network traffic from the client logical partitions is sent over the backup adapter. If and when the primary Shared Ethernet Adapter recovers from its failure, it again begins actively bridging all network traffic.

A Shared Ethernet Adapter in failover mode might optionally have more than one trunk virtual Ethernet. In this case, all the virtual Ethernet adapters in a Shared Ethernet Adapter must have the same priority value. Also, the virtual Ethernet adapter used specifically for the control channel does not need to have the trunk adapter setting enabled. The virtual Ethernet adapters used for the control channel on each Shared Ethernet Adapter in failover mode must have an identical PVID value, and that PVID value must be unique in the system, so that no other virtual Ethernet adapters on the same system are using that PVID.

To ensure prompt recovery times, when you enable the Spanning Tree Protocol on the switch ports connected to the physical adapters of the Shared Ethernet Adapter, you can also enable the portfast option on those ports. The portfast option allows the switch to immediately forward packets on the port without first completing the Spanning Tree Protocol. (Spanning Tree Protocol blocks the port completely until it is finished.)

The Shared Ethernet Adapter is designed to prevent network loops. However, as an additional precaution, you can enable Bridge Protocol Data Unit (BPDU) Guard on the switch ports connected to the physical adapters of the Shared Ethernet Adapter. BPDU Guard detects looped Spanning Tree Protocol BPDU packets and shuts down the port. This helps prevent broadcast storms on the network. A *broadcast storm* is a situation where one message that is broadcast across a network results in multiple responses. Each response generates more responses, causing excessive transmission of broadcast messages. Severe broadcast storms can block all other network traffic, but they can usually be prevented by carefully configuring a network to block disallowed broadcast messages.

Note: When the Shared Ethernet Adapter is using GARP VLAN Registration Protocol (GVRP), it generates BPDU packets, which cause BPDU Guard to shut down the port unnecessarily. Therefore, when the Shared Ethernet Adapter is using GVRP, do not enable BPDU Guard.

For more information about how to enable the Spanning Tree Protocol, the portfast option, and BPDU Guard on the ports, see the documentation that is provided with the switch.

Related tasks

Scenario: Configuring Shared Ethernet Adapter failover
Use this scenario to help you to configure primary and backup **Shared Ethernet Adapters** in the Virtual I/O Server logical partitions.

Shared Ethernet adapters for load sharing

Learn about configuring shared Ethernet adapters (SEA) with load sharing to share the workload between the primary and backup SEA.

The SEA failover configuration provides redundancy only by configuring a backup SEA on a different Virtual I/O Server (VIOS) logical partition. This backup SEA is in the standby mode and can be used only if the primary SEA fails. Hence, the bandwidth of the backup SEA is not used.

On the VIOS Version 2.2.1.0, or later, you can use the SEA failover with load sharing configuration to use the bandwidth of the backup SEA without any impact to reliability.

Note: In load sharing configuration, as with failover, only two SEAs are used. Attempting to configure load sharing with more than two SEAs per vSwitch/PVID is not supported. Additionally, the VLAN IDs on a pair of Shared Ethernet Adapters must not be on any other Shared Ethernet Adapters on the vSwitch.

In the SEA failover with load sharing configuration, the primary and the backup SEAs negotiate those set of virtual local area network (VLAN) IDs, which they are responsible for bridging. After successful negotiation, each SEA bridges the assigned trunk adapters and the associated VLANs. Thus, the primary and the backup SEA bridge the workload for their respective VLANs. If a failure occurs, the active SEA bridges all trunk adapters and the associated VLANs. This action helps to avoid disruption in network services. When the failure is resolved, an SEA automatically returns to the *load sharing* state. Load sharing can also be restarted by running the **chdev** command on the backup SEA. For more information, see chdev command.

To configure SEA failover with load sharing, you must have two or more trunk adapters with distinct VLAN definitions assigned to each SEA. To make optimum use of the SEA failover with load sharing configuration, design the workload such that it is equally distributed among trunk adapters.

Note: When SEA load sharing is configured with Link Aggregation Control Protocol (LACP) (8023ad link aggregation) or physical adapters, the **adapter_reset** value must be set to *no* on both the primary and backup SEA in the VIOS version 2.2.4.0, or earlier, to avoid temporary network outage that might be caused due to a delay in LACP negotiation and a physical adapter reset.

Virtual I/O Server logical partition

Redundancy options for the Virtual I/O Server include multipathing, Redundant Array of Independent Disks (RAID) configurations, and Link Aggregation (or Etherchannel).

Multipathing

Multipathing for the physical storage within the Virtual I/O Server provides failover physical path redundancy and load balancing. The multipathing solutions available in the Virtual I/O Server include MPIO as well as solutions provided by the storage vendors.

For more information about supported storage and multipathing software solutions, see the data sheet available on the Fix Central website.

RAID

Redundant Array of Independent Disks (RAID) solutions provide for device level redundancy within the Virtual I/O Server. Some RAID options, such as LVM mirroring and striping, are provided by the Virtual I/O Server software, while other RAID options are made available by the physical storage subsystem.

See the Virtual I/O Server data sheet available on the Fix Central website for supported hardware RAID solutions.

Link aggregation or Etherchannel devices

A link aggregation, or Etherchannel device, is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters that are aggregated can then act as a single Ethernet device. Link aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` adapters can be aggregated to the `ent3` adapter. The system considers these aggregated adapters as one adapter, and all adapters in the link aggregation device are given the same hardware address. Therefore, they are treated by remote systems as if they were one adapter.

Link aggregation can provide increased redundancy because individual links might fail. The link aggregation device can automatically fail over to another adapter in the device to maintain connectivity. For example, if the `ent0` adapter fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. The `ent0` adapter automatically returns to service on the link aggregation device when it recovers.

You can configure a Shared Ethernet Adapter to use a link aggregation, or Etherchannel, device as the physical adapter.

Redundancy configuration using virtual Fibre Channel adapters

Redundancy configurations help protect your network from physical adapter failures as well as Virtual I/O Server failures.

With N_Port ID Virtualization (NPIV), you can configure the managed system so that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter. Each virtual Fibre Channel adapter is identified by a unique worldwide port name (WWPN), which means that you can connect each virtual Fibre Channel adapter to independent physical storage on a SAN.

Similar to virtual Small Computer Serial Interface (SCSI) redundancy, virtual Fibre Channel redundancy can be achieved by using Multi-path I/O (MPIO) and mirroring at the client partition. The difference between traditional redundancy with SCSI adapters and the NPIV technology by using virtual Fibre Channel adapters, is that the redundancy occurs on the client because only the client recognizes the disk. The Virtual I/O Server is just a pipe. Example 2 uses multiple Virtual I/O Server logical partitions to add redundancy at the Virtual I/O Server level as well.

Example 1: Host bus adapter failover

This example uses Host bus adapter (HBA) failover to provide a basic level of redundancy for the client logical partition. The figure shows the following connections:

- The storage area network (SAN) connects physical storage to two physical Fibre Channel adapters that are located on the managed system.
- The physical Fibre Channel adapters are assigned to the Virtual I/O Server and support NPIV.
- The physical Fibre Channel ports are each connected to a virtual Fibre Channel adapter on the Virtual I/O Server. The two virtual Fibre Channel adapters on the Virtual I/O Server are connected to ports on two different physical Fibre Channel adapters to provide redundancy for the physical adapters.
- Each virtual Fibre Channel adapter on the Virtual I/O Server is connected to one virtual Fibre Channel adapter on a client logical partition. Each virtual Fibre Channel adapter on each client logical partition receives a pair of unique WWPNs. The client logical partition uses one WWPN to log in to the SAN at any given time. The other WWPN is used when you move the client logical partition to another managed system.

The virtual Fibre Channel adapters always have a one-to-one relationship between the client logical partitions and the virtual Fibre Channel adapters on the Virtual I/O Server logical partition. That is, each virtual Fibre Channel adapter that is assigned to a client logical partition must connect to only one virtual Fibre Channel adapter on the Virtual I/O Server, and each virtual Fibre Channel on the Virtual I/O Server must connect to only one virtual Fibre Channel adapter on a client logical partition.

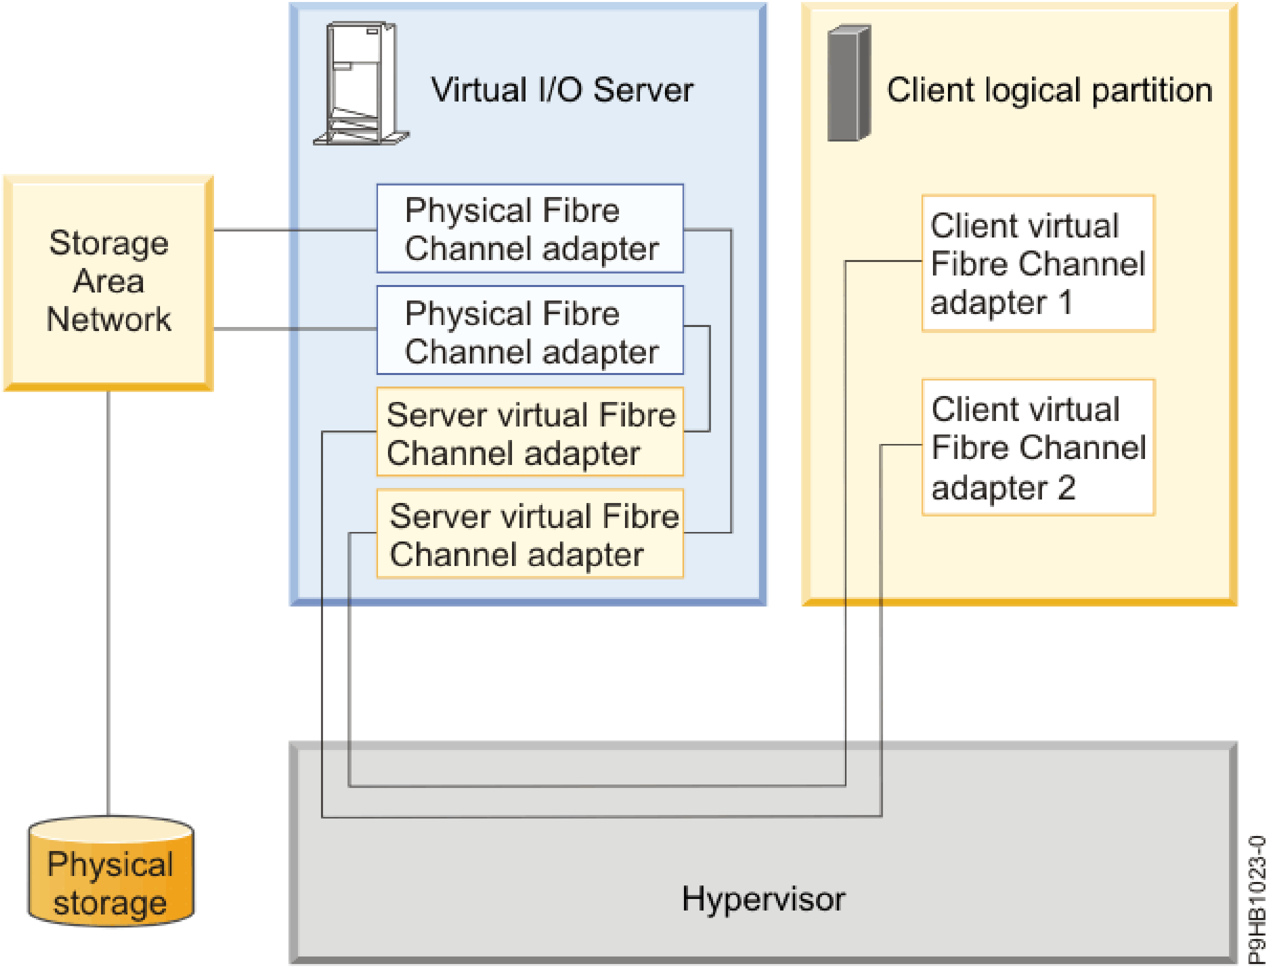
Virtual I/O Server
Client logical partition
Storage Area Network
Physical Fibre Channel adapter
Physical Fibre Channel adapter
Server virtual Fibre Channel adapter
Server virtual Fibre Channel adapter
Client virtual Fibre Channel adapter 1
Client virtual Fibre Channel adapter 2
Physical storage
Hypervisor
P9HB1023-0

The client can write to the physical storage through client virtual Fibre Channel adapter 1 or 2. If a physical Fibre Channel adapter fails, the client uses the alternative path. This example does not show redundancy in the physical storage, but rather assumes it would be built into the SAN.

Note: It is suggested that you configure virtual Fibre Channel adapters from multiple logical partitions to the same HBA, or you configure virtual Fibre Channel adapters from the same logical partition to different HBAs.

Example 2: HBA and Virtual I/O Server failover

This example uses HBA and Virtual I/O Server failover to provide a more advanced level of redundancy for the client logical partition. The figure shows the following connections:

- The storage area network (SAN) connects physical storage to two physical Fibre Channel adapters that are located on the managed system.
- There are two Virtual I/O Server logical partitions to provide redundancy at the Virtual I/O Server level.
- The physical Fibre Channel adapters are assigned to their respective Virtual I/O Server and support NPIV.
- The physical Fibre Channel ports are each connected to a virtual Fibre Channel adapter on the Virtual I/O Server. The two virtual Fibre Channel adapters on the Virtual I/O Server are connected to ports on two different physical Fibre Channel adapters to provide redundancy for the physical adapters. A single adapter might have multiple ports.
- Each virtual Fibre Channel adapter on the Virtual I/O Server is connected to one virtual Fibre Channel adapter on a client logical partition. Each virtual Fibre Channel adapter on each client logical partition receives a pair of unique WWPNs. The client logical partition uses one WWPN to log in to the SAN at any given time. The other WWPN is used when you move the client logical partition to another managed system.

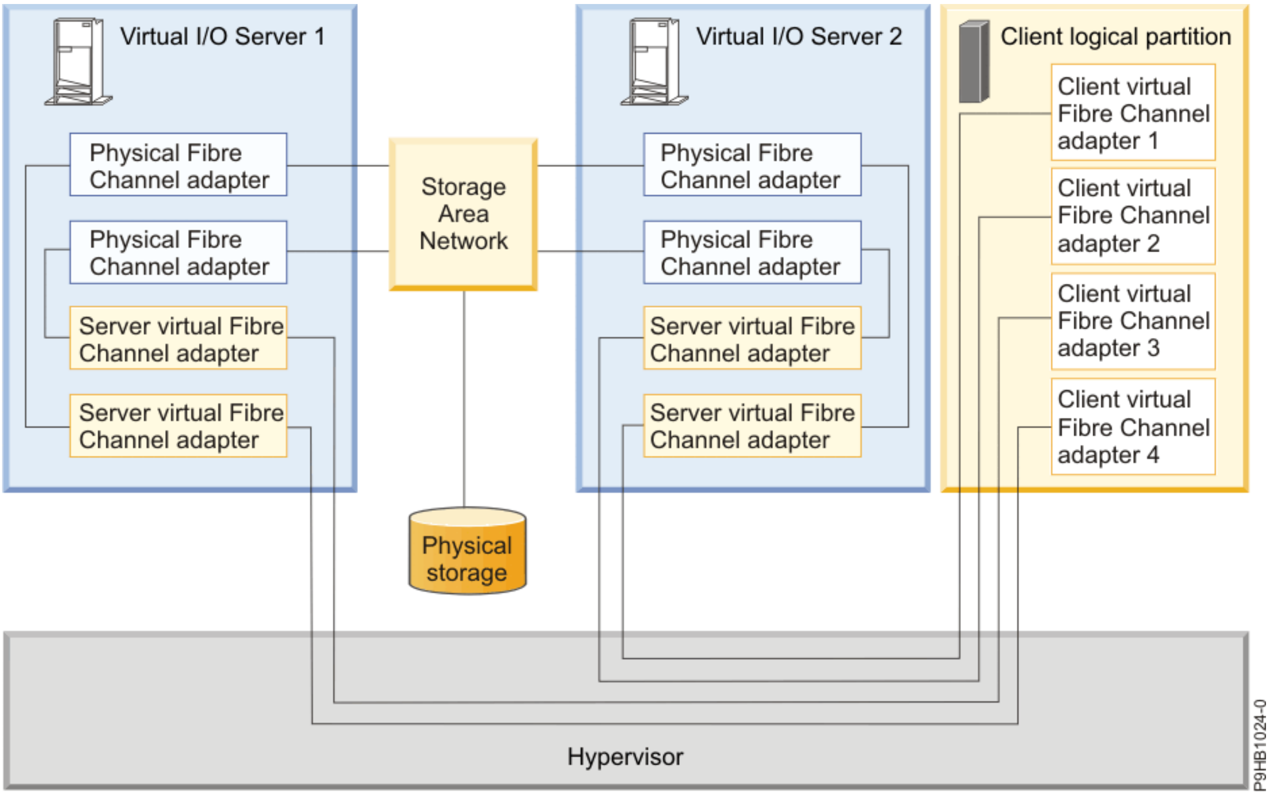
Virtual I/O Server 1
Physical Fibre Channel adapter
Physical Fibre Channel adapter
Server virtual Fibre Channel adapter
Server virtual Fibre Channel adapter
Storage Area Network
Physical storage
Virtual I/O Server 2
Physical Fibre Channel adapter
Physical Fibre Channel adapter
Server virtual Fibre Channel adapter
Server virtual Fibre Channel adapter
Client logical partition
Client virtual Fibre Channel adapter 1
Client virtual Fibre Channel adapter 2
Client virtual Fibre Channel adapter 3
Client virtual Fibre Channel adapter 4
Hypervisor
P9HB1024-0

The client can write to the physical storage through virtual Fibre Channel adapter 1 or 2 on the client logical partition through VIOS 2. The client can also write to physical storage through virtual Fibre Channel adapter 3 or 4 on the client logical partition through VIOS 1. If a physical Fibre Channel adapter fails on VIOS 1, the client uses the other physical adapter connected to VIOS 1 or uses the paths connected through VIOS 2. If VIOS 1 fails, then the client uses the path through VIOS 2. This example does not show redundancy in the physical storage, but rather assumes it would be built into the SAN.

Considerations

These examples can become more complex as you add physical storage redundancy and multiple clients, but the concepts remain the same. Consider the following points:

- To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client logical partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client logical partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.
- Consider load balancing when mapping a virtual Fibre Channel adapter on the Virtual I/O Server to a physical port on the physical Fibre Channel adapter.
- Consider what level of redundancy already exists in the SAN to determine whether to configure multiple physical storage units.
- Consider using two Virtual I/O Server logical partitions. Since the Virtual I/O Server is central to communication between logical partitions and the external network, it is important to provide a level of redundancy for the Virtual I/O Server. Multiple Virtual I/O Server logical partitions require more resources as well, so you must plan accordingly.
- NPIV technology is useful when you want to move logical partitions between servers. For example, in active partition mobility, if you use the redundancy configurations as illustrated, in combination with physical adapters, you can stop all the I/O activity through the dedicated, physical adapter and direct all traffic through a virtual Fibre Channel adapter until the logical partition is successfully moved. The dedicated physical adapter would need to be connected to the same storage as the virtual path. Since you cannot migrate a physical adapter, all I/O activity is routed through the virtual path while you move the partition. After the logical partition is moved successfully, you need to set up the dedicated path (on the destination logical partition) if you want to use the same redundancy configuration you had configured on the original logical partition. Then the I/O activity can resume through the dedicated adapter, by using the virtual Fibre Channel adapter as a secondary path.

Related information

Virtual I/O Server Deployment Examples

Configuring a virtual Fibre Channel adapter using the HMC

IBM PowerVM Live Partition Mobility

Security considerations

Review the security considerations for virtual Small Computer Serial Interface (SCSI), virtual Ethernet, and Shared Ethernet Adapter and the additional security options available.

IBM systems allow cross-partition device sharing and communication. Functions such as dynamic LPAR, shared processors, virtual networking, virtual storage, and workload management all require facilities to ensure that system-security requirements are met. Cross-partition and virtualization features are designed to not introduce any security exposure beyond what is implied by the function. For example, a virtual LAN connection would have the same security considerations as a physical network connection. Carefully consider how to use cross-partition virtualization features in high-security environments. Any visibility between logical partitions must be manually created through administrative system-configuration choices.

Using virtual SCSI, the Virtual I/O Server provides storage to client logical partitions. However, instead of SCSI or fiber cable, the connection for this functionality is done by the firmware. The virtual SCSI device drivers of the Virtual I/O Server and the firmware ensure that only the system administrator of the Virtual I/O Server has control over which logical partitions can access data on Virtual I/O Server storage devices. For example, a client logical partition that has access to a logical volume `lv001` exported by the Virtual I/O Server logical partition cannot access `lv002`, even if it is in the same volume group.

Similar to virtual SCSI, the firmware also provides the connection between logical partitions when using virtual Ethernet. The firmware provides the Ethernet switch functionality. The connection to the external network is provided by the Shared Ethernet Adapter function on the Virtual I/O Server. This part of the Virtual I/O Server acts as a layer-2 bridge to the physical adapters. A VLAN ID tag is inserted into every Ethernet frame. The Ethernet switch restricts the frames to the ports that are authorized to receive frames with that VLAN ID. Every port on an Ethernet switch can be configured to be a member of several VLANs. Only the network adapters, both virtual and physical that are connected to a port (virtual or physical) that belongs to the same VLAN can receive the frames. The implementation of this VLAN standard ensures that the logical partitions cannot access restricted data.

Limitations and restrictions for IBM i client logical partitions

With Virtual I/O Server, you can install IBM i in a client logical partition on POWER8 or POWER9 systems. IBM i client logical partitions have unique system and storage requirements and considerations.

The following limitations and restrictions apply to IBM i client logical partitions of the Virtual I/O Server that are running on HMC-managed systems.

Hardware and software prerequisites

For more information about the supported operating systems, see System software maps.

I/O, storage, and networking limitations for virtual Small Computer Serial Interface (SCSI) adapters

- The IBM i 7.1 TR8, or later client logical partitions can have up to 32 disk units (logical volumes, physical volumes, or files) and up to 16 optical units under a single virtual adapter.
- The maximum virtual disk size is 2 TB minus 512 bytes. If you are limited to one adapter and you have a storage requirement of 32 TB, for example, you might need to make your virtual disks the maximum size of 2 TB. However, in general, consider spreading the storage over multiple virtual disks with smaller capacities. This can help improve concurrency.
- Mirroring and multipath through up to 8 Virtual I/O Server partitions is the redundancy option for client logical partitions. However, you also can use multipathing and RAID on the Virtual I/O Server for redundancy.
- It is required that you assign the tape device to its own Virtual I/O Server adapter, as tape devices often send large amounts of data, which might affect the performance of any other device on the adapter.

SAS adapter performance considerations

If you are using the Virtual I/O Server with Peripheral Component Interconnect Express (PCIe) attached serial-attached SCSI (SAS) adapters to virtualize storage with the IBM i operating system, be aware of specific configuration options that maximize performance. Failure to implement these options can cause write performance degradation. Planning for these considerations ensures that the system is sized for the number of IBM i client logical partitions. For more information about how to configure your Virtual I/O Server, see the SAS Adapter Performance Boost with VIOS topic in the IBM developerWorks website (https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM i Technology Updates/page/SAS Adapter Performance Boost with VIOS).

Virtual Fibre Channel limitations

- The IBM i client partition supports up to 128 target port connections per virtual Fibre Channel adapter.
- The IBM i 7.2 TR7 and IBM i 7.3 TR3 client partitions support up to 127 SCSI devices per virtual Fibre Channel adapter. The 127 SCSI devices can be any combination of disk units or tape libraries. With tape libraries, each control path is counted as a unique SCSI device in addition to a single SCSI device per tape drive.
- For IBM i client partitions, the LUNs of the physical storage connected with NPIV require a storage-specific device driver and do not use the generic virtual SCSI device driver.
- The IBM i client partition supports up to eight multipath connections to a single Fibre Channel disk unit. Each multipath connection can be made with a virtual Fibre Channel adapter or with Fibre Channel I/O adapter hardware that is assigned to the IBM i partition.

- IBM i supports mapping the same physical Fibre Channel port to multiple virtual Fibre Channel adapters in the same IBM i client. All LUNs (disk or tape) that are associated to that physical Fibre Channel adapter must be unique so that there are no multi-path created within the same physical port. In order to use Live Partition Mobility (LPM) or remote restart capability, you can only map the physical port twice to the same IBM i logical partition. The VIOS must be at the version 3.1.2.0, or later and the HMC must be at a version 9.2.950, or later are required for the LPM and restart the logical partition with double mapped ports support.

Installing the Virtual I/O Server and client logical partitions

Find instructions for installing the Virtual I/O Server and client logical partitions by deploying a system plan or manually creating the logical partition and logical partition profiles and installing the Virtual I/O Server (VIOS) and client operating systems.

These instructions apply to installing the Virtual I/O Server and client logical partitions on a system that is managed by a Hardware Management Console (HMC).

The installation procedures vary depending on the following factors:

- The version of HMC attached to the managed system on which you plan to install the Virtual I/O Server and client logical partitions. HMC Version 7, or later displays a different interface than prior versions of the HMC. HMC Version 7, or later also provides the ability to deploy a system plan that includes the Virtual I/O Server and client logical partitions.
- Whether you plan to deploy a system plan that includes the Virtual I/O Server and client logical partitions. When you deploy a system plan, the HMC automatically performs the following tasks based on the information provided in the system plan:
  - Creates the Virtual I/O Server logical partition and logical partition profile.
  - Installs the Virtual I/O Server and provisions virtual resources.
  - Creates the client logical partitions and logical partition profiles.
  - Installs the AIX and Linux operating systems on client logical partitions. The HMC must be at V7R3.3.0, or later.

Related information

Installing the Virtual I/O Server using NIM

Installing the Virtual I/O Server manually by using HMC Version 7 Release 7.1, and later

You can create the Virtual I/O Server logical partition and logical partition profile, and you can install the Virtual I/O Server (VIOS) by using the Hardware Management Console (HMC) Version 7 Release 7.1, or later.

Before you begin

Before you start, ensure that you meet the following requirements:

- The system on which you plan to install the Virtual I/O Server is managed by a Hardware Management Console (HMC).
- The HMC is at Version 7 Release 7.1, or later.

Entering the activation code for PowerVM Editions by using HMC Version 7, or later

Use these instructions to enter the PowerVM Editions activation code by using the Hardware Management Console (HMC) Version 7, or later.

If PowerVM Editions is not enabled on your system, you can use the HMC to enter the activation code that you received when you ordered the feature.

Use the following procedure to enter the activation code for the PowerVM Standard Edition and the PowerVM Enterprise Edition. For more information about the PowerVM Editions, see Introduction to PowerVM.

When the HMC is at version 8.7.0, or later, complete the following steps to enter your activation code:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Capacity on Demand** > **Licensed Capabilities**. The **Licensed Capabilities** page opens.
5. Click **Enter Activation Code**.
6. Enter your activation code and click **OK**.

Creating the Virtual I/O Server logical partition on an HMC managed system

You can use the Hardware Management Console (HMC) Version 7, release 7.1, or later to create a logical partition and partition profile for the Virtual I/O Server (VIOS).

About this task

You can use the Hardware Management Console (HMC) Version 7, release 7.1, or later to create the Virtual I/O Server partition and profile manually. Or, you can deploy a system plan to create the Virtual I/O Server (VIOS) partition and profile. When you deploy a system plan you can optionally create client logical partitions and their profiles on the managed system as well.

For more information about creating a logical partition when the HMC is at version 8.7.0, or later, see Adding a Virtual I/O Server.

For more information about deploying a system plan to create the VIOS when the HMC is at version 8.7.0, or later, see Deploying a system plan by using the HMC.

Creating the Virtual I/O Server logical partition and partition profile manually by using the HMC

You can use the Hardware Management Console (HMC) Version 7, release 7.1, or later to create a logical partition and partition profile for the Virtual I/O Server (VIOS).

Before you begin

Before you start, ensure that you meet the following requirements:

- You are a super administrator or an operator.
- The PowerVM Editions feature is activated. For more information, see “Entering the activation code for PowerVM Editions by using HMC Version 7, or later” on page 90.

About this task

The Virtual I/O Server requires a minimum of 30 GB of disk space.

For more information about creating a logical partition when the HMC is at version 8.7.0, or later, see Creating logical partitions.

What to do next

After you create the partition and partition profile, you are ready to install the Virtual I/O Server. For instructions, see one of the following procedures:

- “Installing the Virtual I/O Server from the HMC command line” on page 92

- “Installing the Virtual I/O Server by using the HMC graphical user interface” on page 92

For more information about adding a Virtual I/O Server when the HMC is at version 8.7.0, or later, see Adding a Virtual I/O Server.

Installing the Virtual I/O Server by using the HMC graphical user interface

You can install the Virtual I/O Server (VIOS) from a CD device, DVD device, saved image, or Network Installation Management (NIM) server by using the Hardware Management Console (HMC) graphical user interface.

For more information about activating and installing the Virtual I/O Server (VIOS) when the HMC is at version 8.7.0, or later, see Activating Virtual I/O Servers.

Installing the Virtual I/O Server from the HMC command line

Find instructions for installing the Virtual I/O Server (VIOS) from the HMC command line by using the **installios** command.

Before you begin

Before you start, complete the following tasks:

1. Ensure that you meet the following requirements:
   - There is an HMC attached to the managed system.
   - The Virtual I/O Server logical partition and logical partition profile are created. For instructions, see “Creating the Virtual I/O Server logical partition and partition profile manually by using the HMC ” on page 91.
   - If you are installing Virtual I/O Server Version 2.2.1.0, or later, ensure that the HMC is at Version 7 Release 7.4.0, or later.
   - The Virtual I/O Server logical partition has at least one Ethernet adapter and a 16 GB disk assigned to it.
   - You have **hmcsuperadmin** authority.
2. Gather the following information:
   - Static IP address for the Virtual I/O Server
   - Subnet mask for the Virtual I/O Server
   - Default gateway for the Virtual I/O Server

About this task

To install the Virtual I/O Server, follow these steps:

Procedure

1. Insert the Virtual I/O Server CD or DVD into the HMC.
2. If you are installing the Virtual I/O Server through the public network interface, continue to step 3. If you are installing the Virtual I/O Server through a private network interface, type the following from the HMC command line:

   ```
   export INSTALLIOS_PRIVATE_IF=interface
   ```

   where, *interface* is the network interface through which the installation must take place.
3. From the HMC command line, type:

   ```
   installios
   ```

4. Follow the installation instructions according to the system prompts.

What to do next

After you install the Virtual I/O Server, finish the installation by checking for updates, setting up remote connections, creating additional user IDs, and so on. For instructions, see "Finishing the Virtual I/O Server installation" on page 93.

Finishing the Virtual I/O Server installation

After you install Virtual I/O Server, you must check for updates, set up remote connects, create additional user IDs, and so on.

Before you begin

This procedure assumes that Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 90.

About this task

To finish the installation, complete the following steps:

Procedure

1. Accept the software maintenance terms and conditions, and the Virtual I/O Server product license. For instructions, see "Viewing and accepting the Virtual I/O Server license" on page 93.
2. Check for updates to the Virtual I/O Server.

   For instructions, see "Updating the Virtual I/O Server" on page 205.
3. Set up remote connections to the Virtual I/O Server.

   For instructions, see "Connecting to the Virtual I/O Server by using OpenSSH" on page 235.
4. Optional: Create the following additional user IDs. After the installation, the only active user ID is the prime administrator (padmin). You can create the following additional user IDs: system administrator, service representative, and development engineer.

   For information about creating user IDs, see "Managing users on the Virtual I/O Server" on page 250.
5. Configure the TCP/IP connection for the Virtual I/O Server using the **mktcpip** command.

   You must complete this task before you can perform any dynamic logical partitioning operations. Alternatively, you can use the configuration assistance menu to configure TCP/IP connections. You can access the configuration assistance menu by running the **cfgassist** command.

What to do next

When you are finished, do one of the following tasks:

- Create client logical partitions.

  Note: You do not need to perform this task if you deployed a system plan to create all your client logical partitions.
- Configure the Virtual I/O Server and install client operating systems. For information, see "Configuring the Virtual I/O Server" on page 109 and Logical partitioning. For more information about Logical partitioning, see Logical partitioning.

Viewing and accepting the Virtual I/O Server license

You must view and accept the license before you use the Virtual I/O Server.

Before you begin

Before you start, ensure that the Virtual I/O Server logical partition profile is created and the Virtual I/O Server is installed. For instructions, see "Installing the Virtual I/O Server and client logical partitions" on page 90.

About this task

To view and accept the Virtual I/O Server license, complete the following steps:

Procedure

1. Log in to the Virtual I/O Server by using the **padmin** user ID.
2. Choose a new password.

   The software maintenance terms and conditions are displayed.
3. If Virtual I/O Server is at Version 1.5 or later, view and accept the software maintenance terms and conditions.

   a) To view the software maintenance terms and conditions, type v on the command line and press enter.

   b) To accept the software maintenance terms and conditions, type a on the command line and press enter.
4. View and accept the Virtual I/O Server product license.

   Note: If you installed the Virtual I/O Server by deploying a system plan, then you already accepted the Virtual I/O Server product license and do not need to complete this step.

   a) To view the Virtual I/O Server product license, type `license -ls` on the command line.

      By default, the license is displayed in English. To change the language in which the license is displayed, follow these steps:

      i) View the list of available locales to display the license by typing the following command:

      ```
      license -ls
      ```

      ii) View the license in another language by typing the following command:

      ```
      license -view -lang Name
      ```

      For example, to view the license in Japanese, type the following command:

      ```
      license -view -lang ja_JP
      ```

   b) To accept the Virtual I/O Server product license, type `license -accept` on the command line.
5. In the installation program, English is the default language. To change the language setting for the system, follow these steps:

   a. View the available languages by typing the following command:

   ```
   chlang -ls
   ```

   b. Change the language by typing the following command, replacing Name with the name of the language you are switching to, as follows:

   ```
   chlang -lang Name
   ```

   Note: If the language file set is not installed, use the `-dev Media` flag to install it.

   For example, to install and change the language to Japanese, type the following command:

   ```
   chlang -lang ja_JP -dev /dev/cd0
   ```

Reinstalling the Virtual I/O Server of a paging VIOS partition

When you reinstall the Virtual I/O Server (VIOS) that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*), you need to reconfigure the shared memory environment. For example, you might need to add the paging space devices again to the shared memory pool.

About this task

The paging VIOS partitions store information about the paging space devices that are assigned to a shared memory pool. The Hardware Management Console (HMC) obtains information about the paging space devices that are assigned to the shared memory pool from the paging VIOS partitions. When you reinstall the VIOS, the information about the paging space devices is lost. For the paging VIOS partitions to regain the information, you must assign the paging space devices again to the share memory pool after you reinstall the VIOS.

The following table shows the reconfiguration tasks that you must perform in the shared memory environment when you resinstall the Virtual I/O Server of a paging VIOS partition.

Table 37. Shared memory reconfiguration tasks for reinstalling the Virtual I/O Server of a paging VIOS partition

Number of paging VIOS partitions that are assigned to the shared memory pool	**Number of paging VIOS partitions for which you want to reinstall the VIOS**	**Reconfiguration steps**	**Instructions**
1	1	1. Shut down all logical partitions that use shared memory (hereafter referred to as *shared memory partitions*). 2. Reinstall the VIOS. 3. Add the paging space devices again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Installing the Virtual I/O Server manually 3. Adding and removing paging space devices to and from the shared memory pool
2	1	1. Shut down each shared memory partition that uses the paging VIOS partition (that you plan to reinstall) as the primary or secondary paging VIOS partition. 2. Remove the paging VIOS partition from the shared memory pool. 3. Reinstall the VIOS. 4. Add the paging VIOS partition again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Removing a paging VIOS partition from the shared memory pool 3. Installing the Virtual I/O Server manually 4. Adding a paging VIOS partition to the shared memory pool

Table 37. Shared memory reconfiguration tasks for reinstalling the Virtual I/O Server of a paging VIOS partition (continued)

Number of paging VIOS partitions that are assigned to the shared memory pool	Number of paging VIOS partitions for which you want to reinstall the VIOS	Reconfiguration steps	Instructions
2	2	1. Shut down all the shared memory partitions. 2. Reinstall the VIOS of each paging VIOS partition. 3. Add the paging space devices again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Installing the Virtual I/O Server manually 3. Adding and removing paging space devices to and from the shared memory pool

Migrating the Virtual I/O Server

You can migrate the Virtual I/O Server (VIOS) logical partition from the Hardware Management Console (HMC) Version 7, or later, from a DVD device that is attached to the Virtual I/O Server logical partition.

Before you start, verify that the following statements are true:

- The system on which you plan to migrate the Virtual I/O Server is managed by a Hardware Management Console (HMC) Version 7, or later.
- The Virtual I/O Server is at Version 1.3, or later.
- The rootvg volume group has been assigned to the Virtual I/O Server.

In most cases, user configuration files from the previous version of the Virtual I/O Server are saved when the new version is installed. If you have two or more Virtual I/O Server logical partitions in your environment for redundancy, you are able to shut down and migrate one Virtual I/O Server logical partition without interrupting any clients. After the migration is complete and the Virtual I/O Server logical partition is running again, the logical partition will be available to clients without additional configuration.

⚠ **Attention:** Do not use the Virtual I/O Server **updateios** command to migrate the Virtual I/O Server.

Related information

Migrating the Virtual I/O Server using NIM

Migrating the Virtual I/O Server from the HMC

Find instructions for migrating the Virtual I/O Server (VIOS) to Version 2.1.0.0, or later, from the Hardware Management Console (HMC) by using the **`installios`** command.

Before you begin

Before you start, verify that you meet the following requirements:

- HMC is attached to the managed system.
- The Virtual I/O Server logical partition has at least one Ethernet adapter and a 16 GB disk assigned to it.
- You have **hmcsuperadmin** authority.
- You have the Virtual I/O Server migration media.

  Note: The migration media is separate from the installation media.
- The Virtual I/O Server is at Version 1.3, or later.

- The disk name (**PV_name**) of your root volume group (rootvg) is `hdisk0`. You can verify the disk name by running the following command from the Virtual I/O Server command line interface: `lsvg -pv rootvg`

  Note: If the disk name is anything other than `hdisk0`, you cannot use the migration DVD to perform the migration. Instead, see Migrating the Virtual I/O Server from a downloaded migration image to ensure that you can migrate the Virtual I/O Server successfully.
- The rootvg volume group has been assigned to the Virtual I/O Server
- Use the **startnetsvc** command to note what services you have started for the Virtual I/O Server.
- Determine the services and agents that are configured (by using the **cfgsvc** command) for use with the Virtual I/O Server. Use the **lssvc** command to display a list of all agents. Use the **lssvc** with the agent name parameter (`lssvc <agent_name>` to display information for a specified agent.

  Note: If any parameters have been set for an agent or service, you will need to reconfigure the parameters after you complete the migration process.
- Back up the mksysb image before migrating Virtual I/O Server. Run the **backupios** command and save the mksysb image to a safe location.

About this task

To migrate the Virtual I/O Server, follow these steps:

Procedure

1. Insert the **Virtual I/O Server migration DVD** into the HMC.
2. If you are installing the Virtual I/O Server through the public network interface, continue to step 3. If you are installing the Virtual I/O Server through a private network interface, type the following command from the HMC command line:

   ```
   export INSTALLIOS_PRIVATE_IF=interface
   ```

   where, *interface* is the network interface through which the installation must take place.
3. From the HMC command line, type:

   ```
   installios
   ```

   ⚠ **Attention:** Do not use the Virtual I/O Server **updateios** command to migrate the Virtual I/O Server.
4. Follow the installation instructions according to the system prompts.

What to do next

After the migration is complete, the Virtual I/O Server logical partition is restarted to its preserved configuration before the migration installation. It is suggested to perform the following tasks:

- Verify that migration was successful by checking the results of the **installp** command and by running the **ioslevel** command. The results of the **ioslevel** command indicate that the ioslevel is now *$ ioslevel 2.1.0.0*.
- Restart previously running daemons and agents:
  1. Log on to the Virtual I/O Server as padmin user.
  2. Complete the following command: `$ motd -overwrite` "*<enter previous banner message>*"
  3. Start up any previously running daemons, such as FTP and Telnet.
  4. Start up any previously running agents, such as ituam.
- Check for updates to the Virtual I/O Server. For instructions, see the Fix Central website.

  Remember: The Virtual I/O Server migration media is separate from the Virtual I/O Server installation media. Do not use the installation media for updates after you perform a migration. It does not contain updates and you will lose your current configuration. Only apply updates by using the instructions from the Virtual I/O Server Support for Power Systems website.

Related tasks

Backing up the Virtual I/O Server to a remote file system by creating an mksysb image
You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to a remote file system by creating an mksysb file.

Migrating the Virtual I/O Server from a downloaded image

Find instructions for migrating the Virtual I/O Server (VIOS) to Version 2.1.0.0, or later, from the Hardware Management Console (HMC) when the disk name of the root volume group (rootvg) is not `hdisk0`.

Before you begin

Ensure that you have the latest HMC installation image. You can obtain the latest installation image from the Fix Central website.

About this task

If the disk name (**`PV_name`**) of your root volume group (rootvg) is anything other than `hdisk0`, complete the following steps to migrate the Virtual I/O Server:

Procedure

1. If the system detects that the first migratable disk does not contain a Virtual I/O Server installation during a non-prompted migration, the migration switches to the prompted mode. At this point, the migration is canceled and the **Migration Confirmation Menu** on the console for the logical partition is displayed with the following message in the screen: `Cannot proceed with VIOS migration. The selected disk does not contain a VIOS.`

   To resolve this problem, you must end the installation process by pressing CTRL-C from the session that executed the `installios` command.
2. Download the Virtual I/O Server migration image from the Virtual I/O Server website.
3. Determine the PVID value for the hard disk of your root volume group (rootvg). There are two ways to obtain the PVID value:

   - From the HMC command line, run the following command: `viosvrcmd -m cecl -p viosl -c "lspv"`

     The command returns information such as in the following example:

     ```
     NAME            PVID                 VG              STATUS
     hdisk0          00cd1b0ef5e5g5g8     None
     hdisk1          00cd1b0ec1b17302     rootvg          active
     hdisk2          none                 None
     ```
   - From the Virtual I/O Server command line with padmin user authority, run the `lspv` to obtain the PVID value of the disk targeted for the installation.

     The command returns information such as in the following example:

     ```
     NAME            PVID                VG              STATUS
     hdisk0          00cd1b0ef5e5g5g8    None
     hdisk1          00cd1b0ec1b17302    rootvg          active
     hdisk2          none                None
     ```
4. From the HMC command line, run the `installios` command with flags. Specify option -E with the PVID value of the Virtual I/O Server target disk that is the migration destination.

For example, based on the following example information, you might run this command: `installios -s cec1 -S 255.255.255.0 -p vios -r vios_prof -i 10.10.1.69 -d /dev/cdrom -m 0e:f0:c0:00:40:02 -g 10.10.1.169 -P auto -D auto -E 00cd1b0ec1b17302`

```
VIOS image source           = /dev/cdrom
managed_system              = cec1
VIOS partition              = vios
VIOS partition profile      = vios_prof
VIOS IP address             = 10.10.1.69
VIOS subnet mask            = 255.255.255.0
VIOS gateway address        = 10.10.1.169
VIOS network MAC address    = 0ef0c0004002
VIOS network adapter speed  = auto
VIOS network adapter duplex = auto
VIOS target disk PVID       = 00cd1b0ec1b17302    rootvg
```

Note: When you install the Virtual I/O Server with the `installios` command, if the installation process cannot find the PVID value that you entered with the -E option, the installation proceeds in the prompt mode.

From the HMC terminal that is running the `installios` command, a message of `info=prompting_for_data_at_console` is displayed. The LED code for the partition shows a code of `0c48`. Either run the `mkvterm -m cec1 -p vios` command from the HMC to interact with the virtual console to continue the migration or to rerun the `installios` command with the corrected PVID value. Note that rerunning the `installios` command recopies the image from media to the disk.

What to do next

After the migration is complete, the Virtual I/O Server logical partition is restarted to its preserved configuration before the migration installation. It is suggested to perform the following tasks:

- Verify that migration was successful by checking the results of the **`installp`** command and by running the **`ioslevel`** command. The results of the **`ioslevel`** command indicate that the ioslevel is now *$ ioslevel 2.1.0.0*.
- Restart previously running daemons and agents:
  1. Log on to the Virtual I/O Server as padmin user.
  2. Complete the following command: `$ motd -overwrite` "*<enter previous banner message>*"
  3. Start up any previously running daemons, such as FTP and Telnet.
  4. Start up any previously running agents, such as ituam.
- Check for updates to the Virtual I/O Server. For instructions, see the Fix Central website.

  Remember: The Virtual I/O Server migration media is separate from the Virtual I/O Server installation media. Do not use the installation media for updates after you perform a migration. It does not contain updates and you can lose your current configuration. Only apply updates using the instructions from the Virtual I/O Server Support for Power Systems website.

Migrating the Virtual I/O Server from DVD

Find instructions for migrating the Virtual I/O Server (VIOS) from a DVD device that is attached to the VIOS logical partition.

Before you begin

Before you start, ensure that you meet the following requirements:

- An HMC is attached to the managed system.
- A DVD optical device is assigned to the Virtual I/O Server logical partition.
- The Virtual I/O Server migration installation media is required.

  Note: The Virtual I/O Server migration installation media is separate from the Virtual I/O Server installation media.

- The Virtual I/O Server is at Version 1.3, or later.
- The root volume group (rootvg) has been assigned to the Virtual I/O Server
- Use the **startnetsvc** command to note what services you have started for the Virtual I/O Server.
- Determine the services and agents that are configured (by using the **cfgsvc** command) for use with the Virtual I/O Server. Use the **lssvc** command to display a list of all agents. Use the **lssvc** with the agent name parameter (`lssvc <agent_name>` to display information for a specified agent.

  Note: If any parameters have been set for an agent or service, you will need to reconfigure the parameters after you complete the migration process.
- Back up the mksysb image before you migrate the Virtual I/O Server. Run the **backupios** command and save the mksysb image to a safe location.

For more information about migrating the Virtual I/O Server (VIOS) from a DVD and activating the (VIOS) when the HMC is at version 8.7.0, or later, see Activating Virtual I/O Servers.

What to do next

After the migration is complete, the Virtual I/O Server logical partition is restarted to its preserved configuration before the migration installation. It is suggested that you perform the following tasks:

- Verify that migration was successful by checking the results of the **installp** command and by running the **ioslevel** command. The results of the **ioslevel** command indicate that the ioslevel is now *$ ioslevel 2.1.0.0*.
- Restart previously running daemons and agents:
  1. Log on to the Virtual I/O Server as `padmin` user.
  2. Complete the following command: `$ motd -overwrite` "*<enter previous banner message>*"
  3. Start any previously running daemons, such as FTP and Telnet.
  4. Start any previously running agents, such as ituam.
- Check for updates to the Virtual I/O Server. For instructions, see the Fix Central website.

  Remember: The Virtual I/O Server migration media is separate from the Virtual I/O Server installation media. Do not use the installation media for updates after you perform a migration. It does not contain updates and you will lose your current configuration. Only apply updates using the instructions from the Virtual I/O Server Support for Power Systems website.

Related tasks

Backing up the Virtual I/O Server to a remote file system by creating an mksysb image
You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to a remote file system by creating an mksysb file.

Migrating the Virtual I/O Server by using the viosupgrade command or by using the manual method

Learn how to upgrade Virtual I/O Server (VIOS) from VIOS version 2.2.x.x. to VIOS version 3.1.0.00. If the VIOS belongs to a Shared Storage Pool (SSP) cluster, the minimum supported level for upgrade to version 3.1 is 2.2.4.x. If the current VIOS version that belongs to an SSP is earlier than 2.2.4.0 (for example, VIOS version 2.2.3.x), you must upgrade to currently supported VIOS versions (such as, VIOS version 2.2.5.x or 2.2.6.x) before attempting to upgrade to VIOS version 2.2.6.30, or later.

The upgrade or migration process on the Virtual I/O Server (VIOS) is different from the upgrade or migration process on other operating systems.

The following tasks must be performed for general VIOS upgrade operations:

- Back up the VIOS metadata by using the `viosbr -backup` command.
- Install a version of VIOS from the available VIOS image.
- Restore the VIOS metadata by using the `viosbr -restore` command.

For more information about upgrading the Virtual I/O Server, see Methods of upgrading a Virtual I/O Server.

Methods of upgrading a Virtual I/O Server

Learn about the methods of upgrading or migrating a Virtual I/O Server (VIOS).

VIOS updates at VIOS version 2.2.x.x are managed through the **updateios** command. The **updateios** command supports only VIOS Technology Level (TL) update operations and not upgrade operations between major versions like version 2 to version 3. For more information, see the updateios command. As the **updateios** command supports only Technology Level (TL) updates, you can use one of the following methods to upgrade to VIOS version 3.1:

- The manual upgrade method
- The new **viosupgrade** tool method

Manual upgrade

In the manual upgrade method, you must first manually back up the VIOS metadata by using the `viosbr -backup` command, install the VIOS through NIM or Flash Storage, and then restore the VIOS metadata by using the `viosbr -restore` command. For more information about the manual upgrade methods when the VIOS belongs to an SSP cluster, see Upgrading the Virtual I/O Server - SSP cluster. For more information about the manual upgrade methods when the VIOS does not belong to an SSP cluster, see Upgrading the Virtual I/O Server - non-SSP cluster

The viosupgrade tool

Before VIOS version 3.1, only the manual upgrade method (backup-install-restore) was available for migration between major versions like version 2 to version 3. In this method, it took some effort for users to repeat the manual process across all the Virtual I/O Servers in their data center. Hence, the **viosupgrade** tool has been developed to provide a single interface to manage the entire VIOS upgrade process automatically. The following two variants of the tool are available:

- NIM - **viosupgrade** for NIM users. For more information, see viosupgrade command
- VIOS - **viosupgrade** for non NIM users. For more information, see viosupgrade command

Notes:

- Installations through the **viosupgrade** command are of the type **New and Complete installation**. Any customized configurations that might exist on the currently running system before the installation starts (including the timezone), are not included in the new installation image. You must save and backup any customized configurations before running the **viosupgrade** command and restore them after the installation completes. The **viosbr backup** and **restore** commands handle only the configurations related to the virtual I/O. The **viosupgrade** command provides an option to save the required configuration files from the currently installed image to the new VIOS image.

  For example, to copy any customized configuration files such as `/etc/netsvc.conf`, `/etc/ntp.conf`, and so on, to the new image, use the **viosupgrade** command. For more information, see the viosupgrade command.
- Alternate disks that are used with the **-a** flag and the **-r** flag as part of the **viosupgrade** command must be completely free. That is, you must be able to list them by using the `lspv -free` command on the VIOS.
- When you are using the **viosupgrade** command, do not make any changes to the virtual device mappings on the VIOS. If you create or change any mappings during the **viosupgrade** process, the new mappings are lost.

Upgrading the Virtual I/O Server - non-SSP cluster

Learn about the process of upgrading or migrating the Virtual I/O Server (VIOS), when the VIOS does not belong to a Shared Storage Pool (SSP) cluster.

You can directly migrate the VIOS from VIOS version 2.2.x.x to VIOS version 3.1, if VIOS does not belong to a Shared Storage Pool (SSP) cluster. The different upgrade processes are explained in the following sections.

Using the `viosupgrade` command with the `bosinst` option from NIM Master

Virtual I/O Servers from VIOS version 2.2.x.x to VIOS version 3.1 can be upgraded by using the **viosupgrade** command from the NIM `bosinst` method.

For a VIOS node that is at version 2.2.4.x or 2.2.5.x and that must be upgraded to version 3.1.0.00, the VIOS node can be directly upgraded to version 3.1.0.00, by using the following command:

```
viosupgrade -t bosinst -n <hostname> -m <mksysb_image> -p <spot_name> -a <hdisk>
```

For example: `viosupgrade -t bosinst -n vios1 -m vios_3.1.0.0 -p vios_3.1.0.0_spot -a hdisk1`

You can check the status of the VIOS installation by using the **viosupgrade -q vios1** command.

Using the `viosupgrade` command with the `altdisk` option from NIM Master

Virtual I/O Servers from VIOS version 2.2.6.30 to VIOS version 3.1 can be upgraded by using the **viosupgrade** command from the NIM `altdisk` method.

The VIOS can be upgraded to version 3.1.0.00, by using the following command:

```
viosupgrade -t altdisk -n <hostname> -m mksysb_image> -p <spot_name> -a <hdisk>
```

For example: `viosupgrade -t altdisk -n vios1 -m vios_3.1.0.0 -p vios_3.1.0.0_spot -a hdisk1`

You can check the status of the VIOS installation by using the **viosupgrade -q vios1** command.

Note: The **viosupgrade -t altdisk** command is supported in VIOS version 2.2.6.30, or later. Hence, this option is not applicable for upgrades where the VIOS is at versions earlier than 2.2.6.30.

Using the `viosupgrade` command from VIOS - Non-NIM environment

Virtual I/O Servers from VIOS version 2.2.6.30 to VIOS version 3.1.0.00, or later, can be upgraded by using the **viosupgrade** command. In this way, you can upgrade a VIOS in a non-NIM environment, where it uses the **alt_disk_mksysb** command to install VIOS version 3.1.0.00 on the provided disk.

You can upgrade the VIOS, by using the following command:

```
viosupgrade -l -i <mksysb image> -a <hdisk>
```

For example: `viosupgrade -l -i vios3.1_mksysb -a hdisk1`

You can check the status of the VIOS installation by using the **viosupgrade -l -q** command.

Traditional method - manual

In the traditional method, you must back up the VIOS metadata by using the **viosbr -backup** command, save the backup in a remote location, install the VIOS by using the available VIOS version, copy the backup data back to the VIOS after the installation, and then restore the VIOS metadata by using the **viosbr -restore** command.

To manually back up and restore the VIOS metadata, complete the following steps:

1. Back up the VIOS metadata, by using the following command:

   ```
   viosbr -backup -file <FileName>
   ```

   Note: You must transfer the backup file (FileName) to a remote location so that you can restore the VIOS metadata when you are in step 4. You can also back up any other data from `rootvg`, if required.
2. Install the VIOS image by using the available installation methods, such as NIM installation.
3. Transfer the backup file (FileName), saved in step 1, to the VIOS.

   Note: Network connectivity must be available for the file transfer to succeed.
4. Restore the VIOS metadata, by using the following command:

   ```
   viosbr -restore -file <FileName>
   ```

   Note: Any other data that was backed up in step 1 can be transferred back to the VIOS , if required.

Upgrading the Virtual I/O Server - SSP cluster

Learn about the process of upgrading or migrating the Virtual I/O Server (VIOS), when the VIOS belongs to a Shared Storage Pool (SSP) cluster.

You can use one of the following methods to upgrade Virtual I/O Servers that belong to a Shared Storage Pool (SSP):

- Non-disruptive upgrades
- Disruptive upgrades

Non-disruptive upgrades

A general recommendation is that Virtual I/O configurations must be through dual VIOS environments. This configuration ensures that an alternative path is always available for I/O communication from client logical partitions in case the primary path goes offline. For non-disruptive upgrades, you can start the upgrade of all Virtual I/O Servers in the primary path, while keeping the Virtual I/O Servers in the alternative path active. During the upgrade process, the cluster and exported Logical Units (LUs) remain available to the client logical partitions through VIOS cluster nodes in the alternative path. Client logical partitions can actively read and write to the SSP Logical Units through other available Virtual I/O paths. After you upgrade the primary Virtual I/O Servers and adding them back to the cluster by using the **`viosbr -restore`** command, you can upgrade the Virtual I/O Servers in the alternative path by repeating the same process.

For a VIOS that belongs to an SSP cluster, if you plan for a non-disruptive update to VIOS version 3.1 you must use the following 2-step upgrade process.

1. Upgrade all the SSP nodes from VIOS versions 2.2.4.x, or later to VIOS version 2.2.6.30, or later, where the VIOS version must be equal to, or greater than 2.2.6.30 and less than version 3.1. After you upgrade all the nodes, wait for the rolling upgrade process to complete in the background, where the contents of the old database are migrated to the new database. The rolling upgrade process is purely internal to the SSP cluster and no action is necessary from you to start this process.
2. As a second step, upgrade all the SSP nodes from VIOS version 2.2.6.30, or later to VIOS version 3.1.0.00, or later.

Upgrades from VIOS 2.2.4.x, or later to VIOS 2.2.6.30

You can choose one of the following methods to upgrade from VIOS version 2.2.4.x, or later to VIOS version 2.2.6.30.

- **Using the `updateios` command**:
  – The **`updateios`** command updates the VIOS to the necessary maintenance level. You do not need to take a backup or restore the VIOS metadata, as no new installation takes place.

  - You can upgrade the VIOS to version 2.2.6.30, by using the following command:

    ```
    updateios -dev <update image location>
    ```

    For example: `updateios -dev /home/padmin/update`

- **Using the `viosupgrade` command from NIM Master – `bosinst` method**:

  The **`viosupgrade`** command from NIM Master is not supported in shared storage pool cluster environment for VIOS versions earlier than 2.2.6.30. You can update the VIOS by using the **`updateios`** command or by performing a Manual Backup-Install-Restore.

- **Manual Backup-Install-Restore**:

  For more information about this method, see Traditional method - manual.

Upgrades from VIOS 2.2.6.30, or later to VIOS 3.1x.x

You can choose the following methods to upgrade from VIOS version 2.2.6.30, or later to VIOS version 3.1.0.00, or later. The **`viosupgrade`** command from NIM Master, supports a complete VIOS upgrade process (backup-install-restore).

Note: To upgrade to VIOS version 3.1 from VIOS version 2.2.6.30 or later, the status of the SSP cluster must be **ON_LEVEL** for all the nodes. You can verify the status of the cluster by using the command, `cluster -status -verbose`.

If the status of the SSP cluster is **UP_LEVEL**, your cluster nodes (Virtual I/O Servers) are not ready for migration to VIOS version 3.1.

During the 2-step upgrade process, to upgrade from a VIOS that is at a version earlier than 2.2.6.30 to a VIOS version 3.1, or later, it is mandatory for the cluster to be at **ON_LEVEL** for all the SSP cluster nodes after the first upgrade step takes the cluster nodes to version 2.2.6.30, or later. When the last node in the cluster gets upgraded to level 2.2.6.30, or later, the SSP internal process called **Rolling Upgrade** starts and migrates the contents of the SSP database from the older version to the currently installed version. The **ON_LEVEL** status for all SSP cluster nodes indicates the completion of step 1 of the upgrade process.

- **Using the `viosupgrade -bosinst` command from NIM Master**:
  - You can upgrade the VIOS to version 3.1.0.00, by using the following command:

    ```
    viosupgrade -t bosinst -n <hostname> -m <mksysb_image> -p <spot_name> -a <hdisk> -c
    ```

    For example: `viosupgrade -t bosinst -n vios1 -m vios_3.1.0.0 -p vios_3.1.0.0_spot -a hdisk1 -c`
  - You can check the status of the VIOS installation by using the **`viosupgrade -q vios1`** command.
- **Using the `viosupgrade -altdisk` command from NIM Master**:
  - To avoid downtime during VIOS installations, you can use the *NIM altdisk* method. This method preserves the current rootvg image and installs the VIOS on a new disk by using the `alt_disk_mksysb` method.
  - You can upgrade the VIOS to version 3.1.0.00, by using the following command:

    ```
    viosupgrade -t altdisk -n <hostname> -m mksysb_image> -p <spot_name> -a <hdisk> -c
    ```

    For example: `viosupgrade -t altdisk -n vios1 -m vios_3.1.0.0 -p vios_3.1.0.0_spot -a hdisk1 -c`
  - You can check the status of the VIOS installation by using the **`viosupgrade -q vios1`** command.

    Note: The `viosupgrade -altdisk` option is supported in VIOS version 2.2.6.30, or later. Hence, this option is not applicable to upgrades with VIOS versions earlier than 2.2.6.30.
- **Using the `viosupgrade` command from VIOS – non-NIM environment**:
  - In a non-NIM environment, you can also use the **`viosupgrade`** command from the VIOS to upgrade the VIOS. For this method, you do not need a NIM master. The **`viosupgrade`** command must be

run directly on the VIOS. This method uses the **alt_disk_mksysb** command to install VIOS version 3.1.0.00 on the provided disk.

- You can upgrade the VIOS to version 3.1.0.00, by using the following command:

  ```
  viosupgrade -l -i <mksysb image> -a <hdisk>
  ```

  For example: `viosupgrade -l -i vios3.1_mksysb -a hdisk1`
- You can check the status of the VIOS installation by using the **viosupgrade -l -q** command.

  Note: The `viosupgrade -altdisk` option is supported in VIOS version 2.2.6.30, or later. Hence, this option is not applicable for upgrades with VIOS versions earlier than 2.2.6.30.

- **Manual Backup-Install-Restore**:

  For more information about this method, see Traditional method - manual.
- **Traditional method - manual**:

  In the traditional method, you must back up the clusters by using the **viosbr -backup -cluster** command, save the backup in a remote location, install the VIOS by using the available VIOS version, copy the backup data back to the VIOS after the installation, and then restore the VIOS metadata by using the **viosbr -restore** command.

  To back up the cluster-level VIOS metadata, complete the following steps:

  1. Back up the cluster-level VIOS metadata, by using the following command:

     ```
     viosbr -backup -clustername <clusterName> -file <FileName>
     ```

     Note: Save the file (FileName) in some location and place it on the VIOS after step 2 is complete to restore the VIOS metadata.
  2. Install the VIOS image by using the available installation methods, such as NIM installation.

     Note: If the VIOS is part of a cluster and the Shared Ethernet Adapter (SEA) is configured on the Ethernet interface used for cluster communication, you must restore the network configuration before you restore the cluster. To restore the network configuration before cluster restore, complete step 3. If you encounter any errors during step 3, you can use the -force flag to continue restoring the network configuration. If SEA is not configured on the Ethernet interface used for cluster communication , then directly go to step 4.
  3. Restore all the network configurations before restoring the cluster, by using the following steps::

     ```
     viosbr -restore -file <FileName> -type net
     ```

     Note: The backup file needs to be copied to the VIOS before starting the restore process. Complete the following steps if there is no IP address configured on the VIOS to transfer the backup file.

     a. Temporarily configure the IP address on an Ethernet interface.

     b. Transfer the backup file to the VIOS.

     c. Undo the configuration done in step a.

     d. Login to the VIOS and execute the **viosbr -restore** command as mentioned in step 3.
  4. Restore the cluster, by using the following command:

     ```
     viosbr -restore -clustername <clusterName> -file <FileName> -repopvs <list_of_disks>
     -curnode
     ```

Disruptive upgrades

In an SSP cluster environment, if you choose disruptive updates, the client logical partitions are extremely likely to go offline as the Logical Units (LUs) of the SSP will not be available as the cluster will be in the offline state during upgrades. To perform this type of upgrade, you must handle the upgrade process manually.

The upgrade is disruptive regarding I/O communication in the client logical partitions. As part of disruptive upgrades, the SSP cluster and all the Logical Units (LUs) that belong to the cluster are offline during the upgrade process. If client logical partitions actively use this storage (such as *rootvg*), it would result in I/O failures and it is extremely likely to cause disruption to logical partition services. In case the rootvg disk of the client logical partition is from SSP LUs, it is not advisable to choose disruptive upgrades, unless the downtime of the client logical partition is acceptable.

Note: The **viosupgrade** command does not support disruptive upgrades.

Supported Virtual I/O Server upgrade levels

Learn about the Virtual I/O Server (VIOS) upgrade levels for Shared Storage Pool (SSP) and non-SSP clusters.

The following table provides details about the supported VIOS upgrade levels.

Table 38. Supported VIOS upgrade levels

VIOS type	Current level of VIOS	Target level	Allowed to upgrade?	VIOS upgrade method to be used
VIOS non-SSP node	< 2.2.6.30	2.2.6.30	Yes	Use the **updateios** command or the Manual Backup-Install-Restore method.
	< 2.2.6.30	3.1.0.00	Yes	Use the **viosupgrade -bosinst** command or the Manual Backup-Install-Restore method.
	2.2.6.30	3.1.0.00	Yes	Use the **viosupgrade -bosinst/ altdisk** command or the Manual Backup-Install-Restore method.

Table 38. Supported VIOS upgrade levels (continued)

VIOS type	Current level of VIOS	Target level	Allowed to upgrade?	VIOS upgrade method to be used
VIOS - SSP node - non-Disruptive	< 2.2.6.30	2.2.6.30	Yes	Use the **updateios** command or the Manual Backup-Install-Restore method.
	< 2.2.6.30	3.1.0.00	No	Not supported
	2.2.6.30	3.1.0.00	Yes	Use the **viosupgrade -bosinst/ altdisk** command or the Manual Backup-Install-Restore method.
VIOS - SSP node - Disruptive	< 2.2.6.30	2.2.6.30	Yes	Use the **updateios** command or the Manual Backup-Install-Restore method.
	< 2.2.6.30	3.1.0.00	Yes	Use the Manual Backup-Install-Restore method.
	2.2.6.30	3.1.0.00	Yes	Use the **viosupgrade -bosinst/ altdisk** command or the Manual Backup-Install-Restore method.

Miscellaneous information about upgrading Virtual I/O Server

This topic provides additional information about the Virtual I/O Server (VIOS) upgrade process.

Third-party software

As the VIOS 3.1 release is a new and fresh installation from a VIOS *mksysb* image, any third-party software must be reinstalled after the VIOS installation. If you upgrade from VIOS version 2.2.x, you can back up the VIOS metadata by using the **viosbr** command. While the **viosbr** command can back up the VIOS metadata, it does not directly manage the metadata of third-party software. Therefore, you must save any third-party metadata, including third-party license requirements before you upgrade the VIOS.

Safe installations

You must upgrade each VIOS in a manner that minimizes impact to your environment. You must ensure that redundant VIOS nodes are taken offline one at a time. For example, the second VIOS must not be taken offline until the first VIOS is back online and fully operational. You must also shut down any client systems that have *rootvg* based virtual devices. Failure to gracefully shut down these nodes might lead to a crash of those nodes.

VIOS image preparation

You can use the VIOS upgrade tool to prepare the VIOS image for your environment.

You can use one of the following VIOS images to deploy in your environment:

- VIOS *mksysb* image that is supplied by IBM.
- Customized VIOS image that is prepared at your location based on your requirement.

Typically, you can download the VIOS *mksysb* image from the IBM website and customize it by installing third-party software. Some such software applications that you might want to include with the VIOS image are, Multipath application drivers, Security profiles, Performance monitoring tools, and so on. After you install the necessary applications or drivers, a customized VIOS *mksysb* image is created. Then, you can use the customized VIOS image to deploy on all the Virtual I/O Servers across the data center.

You can create a customized *mksysb* image, by using the following command:

```
backupios -mksysb -file <filename.mksysb>
```

Unsupported upgrade scenarios with the viosupgrade tool

Learn about the scenarios that are not supported for upgrade with the Virtual I/O Server (VIOS) upgrade tool.

Full cluster restore in a single instance is not supported

The VIOS upgrade tool supports backup and restore of VIOS at cluster level with the **-c** flag. However, full cluster restore in a single instance is not supported in the VIOS 3.1 release. Hence, when you upgrade VIOS nodes that belong to a Shared Storage Pool (SSP) cluster, irrespective of the number of nodes, you must upgrade a few nodes at a time, while keeping the cluster running on the other nodes. Failure to do so might result in the loss of cluster connectivity.

For example, for a 4-node cluster, you can upgrade 1, 2, or 3 nodes, while keeping at least one node active in the cluster. After the successful installation of the first set of nodes, you can choose to upgrade the second set of nodes in the cluster.

Note: For single node clusters, it is not possible to use the VIOS upgrade tool to upgrade and restore the cluster. You must manage single node clusters manually. Alternatively, you can add one or more nodes to the SSP cluster before you initiate the upgrade process on the first node.

Rootvg LV backed vSCSI disk backup restore not supported

Currently, the **viosbr** command does not support virtual SCSI (vSCSI) disks that are created on the *rootvg* disks of the Virtual I/O Server. Hence, the VIOS upgrade tool cannot be used to restore vSCSI mappings if Logical Volumes (LVs) are created from a *rootvg* disk. You must move the vSCSI disks from the *rootvg* to other volume groups before you initiate the upgrade process. Alternatively, you can start the installation on an alternative disk by preserving the current *rootvg*.

For more information about using Logical Volume Manager (LVM) commands (such as the **cplv** command) to migrate these vSCSI logical volumes, see IBM Moving a JFS/JFS2 File System to a New Volume Group.

Configuring the Virtual I/O Server

You need to configure virtual Small Computer Serial Interface (SCSI) and virtual Ethernet devices on the Virtual I/O Server. Optionally, you can also configure virtual Fibre Channel adapters, Tivoli agents and clients, and configure the Virtual I/O Server as an LDAP client.

Configuring virtual SCSI on the Virtual I/O Server

You can configure virtual Small Computer Serial Interface (SCSI) devices by deploying a system plan, creating volume groups and logical volumes, and configuring the Virtual I/O Server to support SCSI-2 reserve functions.

About this task

Provisioning virtual disk resources occurs on the Virtual I/O Server. Physical disks that are owned by the Virtual I/O Server can either be exported and assigned to a client logical partition as a whole or can be partitioned into parts, such as logical volumes or files. These logical volumes and files can be exported as virtual disks to one or more client logical partitions. Therefore, by using virtual SCSI, you can share adapters as well as disk devices.

To make a physical volume, logical volume, or file available to a client logical partition requires that it be assigned to a virtual SCSI server adapter on the Virtual I/O Server. The SCSI client adapter is linked to a particular virtual SCSI server adapter in the Virtual I/O Server logical partition. The client logical partition accesses its assigned disks through the virtual SCSI client adapter. The Virtual I/O Server client adapter sees standard SCSI devices and LUNs through this virtual adapter. Assigning disk resources to a SCSI server adapter in the Virtual I/O Server effectively allocates resources to a SCSI client adapter in the client logical partition.

For more information about SCSI devices that you can use, see the Fix Central website.

Creating the virtual target device on the Virtual I/O Server

Creating the virtual target device on the Virtual I/O Server maps the virtual Small Computer Serial Interface (SCSI) adapter with the file, logical volume, tape, optical device, or physical disk.

About this task

With the Virtual I/O Server Version 2.1, and later, you can export the following types of physical devices:

- Virtual SCSI disk that is backed by a physical volume
- Virtual SCSI disk that is backed by a logical volume
- Virtual SCSI disk that is backed by a file
- Virtual SCSI optical backed by a physical optical device
- Virtual SCSI optical backed by a file
- Virtual SCSI tape that is backed by a physical tape device

After a virtual device is assigned to a client partition, the Virtual I/O Server must be available before the client logical partitions can access it.

Creating a virtual target device on a Virtual I/O Server that maps to a physical or logical volume, tape or physical optical device

You can create a virtual target device on a Virtual I/O Server that maps the virtual Small Computer Serial Interface (SCSI) adapter to a physical disk, tape, or physical optical device, or to a logical volume that is based on a volume group.

About this task

The following procedure can be repeated to provide additional virtual disk storage to any client logical partition.

Before you start, ensure that the following statements are true:

1. At least one physical volume, tape, or optical device, or logical volume is defined on the Virtual I/O Server. For more information, see "Logical volumes" on page 17.
2. The virtual adapters for the Virtual I/O Server and the client logical partitions are created. This usually occurs during the creation of the logical partition profile. For more information about creating the logical partition, see Installing the Virtual I/O Server.
3. Be aware of the maximum transfer size limitation when you use AIX clients and physical devices. If you have an existing and active AIX client, and you want to add another virtual target device to the virtual SCSI server adapter used by that client, ensure that the max_transfer attribute is the same size or larger than the devices already in use.
4. Be aware of the maximum transfer size limitation when you use AIX clients and physical devices. If you have an existing and active AIX client, and you want to add another virtual target device to the virtual SCSI server adapter used by that client, ensure that the max_transfer attribute is the same size or larger than the devices already in use.

Tip: If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create a virtual target device on the Virtual I/O Server.

To create a virtual target device that maps a virtual SCSI server adapter to a physical device or logical volume, complete the following steps from the Virtual I/O Server command-line interface:

Procedure

1. Use the **`lsdev`** command to ensure that the virtual SCSI adapter is available. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   ent3     Available  Virtual I/O Ethernet Adapter (l-lan)
   vhost0   Available  Virtual SCSI Server Adapter
   vhost1   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtscsi0  Available  Virtual Target Device - Logical Volume
   vtscsi1  Available  Virtual Target Device - File-backed Disk
   vtscsi2  Available  Virtual Target Device - File-backed Disk
   ```

2. To create a virtual target device, which maps the virtual SCSI server adapter to a physical device or logical volume, run the **`mkvdev`** command:

   ```
   mkvdev -vdev TargetDevice -vadapter VirtualSCSIServerAdapter
   ```

   Where:

   - *TargetDevice* is the name of the target device, as follows:
     - To map a logical volume to the virtual SCSI server adapter, use the name of the logical volume. For example, lv_4G.
     - To map a physical volume to the virtual SCSI server adapter, use `hdisk`*x*. For example, hdisk5.
     - To map an optical device to the virtual SCSI server adapter, use `cd`*x*. For example, cd0.
     - To map a tape device to a virtual SCSI adapter, use `rmt`*x*. For example, rmt1.

- *VirtualSCSIServerAdapter* is the name of the virtual SCSI server adapter.

Note: If needed, use the **lsdev** and **lsmap -all** commands to determine the target device and virtual SCSI server adapter that you want to map to one another.

The storage is available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed (on a Linux logical partition), or configured (on an AIX logical partition), or appears as a either a DDXXX or DPHXXX device (on an IBM i partition).

3. View the newly created virtual target device by running the **lsdev** command.

   For example, running `lsdev -virtual` returns results similar to the following:

```
name     status     description
vhost3   Available  Virtual SCSI Server Adapter
vsa0     Available  LPAR Virtual Serial Adapter
vtscsi0  Available  Virtual Target Device - Logical Volume
vttape0  Available  Virtual Target Device - Tape
```

4. View the logical connection between the newly created devices by running the **lsmap** command.

   For example, running `lsmap -vadapter vhost3` returns results similar to the following:

```
SVSA      Physloc                      Client PartitionID
-------------------------------------------------------
vhost3    U9111.520.10DDEEC-V1-C20   0x00000000

VTD                     vtscsi0
Status                  Available
LUN                     0x8100000000000000
Backing device          lv_4G
Physloc
```

   The physical location is a combination of the slot number, in this case 20, and the logical partition ID. The storage is now available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed, or configured.

What to do next

If you later need to remove the virtual target device, you can do so by using the **rmvdev** command.

Related concepts

Virtual SCSI sizing considerations

Understand the processor and memory-sizing considerations when you implement virtual Small Computer Serial Interface (SCSI).

Related information

Creating a virtual disk for a VIOS logical partition using the HMC

Virtual I/O Server commands

Creating a virtual target device on a Virtual I/O Server that maps to a file or logical volume

You can create a virtual target device on a Virtual I/O Server that maps the virtual Small Computer Serial Interface (SCSI) adapter to a file or a logical volume that is based on a storage pool.

About this task

The following procedure can be repeated to provide additional virtual disk storage to any client logical partition.

Before you start, ensure that the following statements are true:

- The Virtual I/O Server is at Version 1.5, or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.
- At least one file is defined in a file storage pool, or at least one logical volume is defined in a logical volume storage pool on the Virtual I/O Server. For more information, see "Virtual storage" on page 21 and "Storage pools" on page 30.

- The virtual adapters for the Virtual I/O Server and the client logical partitions are created. This usually occurs during the creation of the logical partition profile. For more information about creating the logical partition, see Installing the Virtual I/O Server.

Tip: If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create a virtual target device on the Virtual I/O Server.

To create a virtual target device that maps a virtual SCSI server adapter to a file or logical volume, complete the following steps from the Virtual I/O Server command-line interface:

Procedure

1. Use the **`lsdev`** command to ensure that the virtual SCSI adapter is available. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   ent3     Available  Virtual I/O Ethernet Adapter (l-lan)
   vhost0   Available  Virtual SCSI Server Adapter
   vhost1   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtscsi0  Available  Virtual Target Device - Logical Volume
   vtscsi1  Available  Virtual Target Device - File-backed Disk
   vtscsi2  Available  Virtual Target Device - File-backed Disk
   ```

2. To create a virtual target device, which maps the virtual SCSI server adapter to a file or logical volume, run the **`mkbdsp`** command:

   ```
   mkbdsp -sp StoragePool -bd BackingDevice -vadapter VirtualSCSIServerAdapter -tn
   TargetDeviceName
   ```

   Where:

   - *StoragePool* is the name of the storage pool that contains the file or logical volume to which you plan to map the virtual SCSI server adapter. For example, fbPool.
   - *BackingDevice* is the name of the file or logical volume to which you plan to map the virtual SCSI server adapter. For example, devFile.
   - *VirtualSCSIServerAdapter* is the name of the virtual SCSI server adapter. For example, vhost4.
   - *TargetDeviceName* is the name of the target device. For example, fbvtd1.

   The storage is available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed (on a Linux logical partition), or configured (on an AIX logical partition), or appears as a either a DDXXX or DPHXXX device (on an IBM i logical partition).

3. View the newly created virtual target device by running the **`lsdev`** command.

   For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   vhost4   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   fbvtd1   Available  Virtual Target Device - File-backed Disk
   ```

4. View the logical connection between the newly created devices by running the **`lsmap`** command.

   For example, running `lsmap -vadapter vhost4` returns results similar to the following:

   ```
   SVSA     Physloc                     Client PartitionID
   -------------------------------------------------------
   vhost4   U9117.570.10C8BCE-V6-C2     0x00000000

   VTD              fbvtd1
   Status           Available
   LUN              0x8100000000000000
   Backing device   /var/vio/storagepools/fbPool/devFile
   Physloc
   ```

   The physical location is a combination of the slot number, in this case 2, and the logical partition ID. The virtual device can now be attached from the client logical partition.

What to do next

If you later need to remove the virtual target device and backup device (file or logical volume), use the **rmbdsp** command. An option is available on the **rmbdsp** command to remove the virtual target device without removing the backup device. A backup device file is associated with a virtual target device by inode number rather than by file name, so do not change the inode number of a backing device file. The inode number might change if you alter a backup device file (by using the AIX **rm**, **mv**, and **cp** commands) (by using the AIX **rm**, **mv**, and **cp** commands), while the backup device file is associated with a virtual target device.

Related information

Creating a virtual disk for a VIOS logical partition using the HMC

Virtual I/O Server commands

Creating a virtual target device on a Virtual I/O Server that maps to a file-backed virtual optical device

You can create a virtual target device on a Virtual I/O Server that maps the virtual Small Computer Serial Interface (SCSI) adapter to a file-backed virtual optical device.

About this task

The following procedure can be repeated to provide additional virtual disk storage to any client logical partition.

Before you start, complete the following steps:

1. Ensure that the Virtual I/O Server is at Version 1.5, or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.
2. Ensure that the virtual adapters for the Virtual I/O Server and the client logical partitions are created. This usually occurs during the creation of the logical partition profile. For more information about creating the logical partition, see "Installing the Virtual I/O Server and client logical partitions" on page 90.

Tip: If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create a virtual target device on the Virtual I/O Server.

To create a virtual target device that maps a virtual SCSI server adapter to a file-backed virtual optical device, complete the following steps from the Virtual I/O Server command-line interface:

Procedure

1. Use the **lsdev** command to ensure that the virtual SCSI adapter is available. For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   ent3     Available  Virtual I/O Ethernet Adapter (l-lan)
   vhost0   Available  Virtual SCSI Server Adapter
   vhost1   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtscsi0  Available  Virtual Target Device - Logical Volume
   vtscsi1  Available  Virtual Target Device - File-backed Disk
   vtscsi2  Available  Virtual Target Device - File-backed Disk
   ```

2. To create a virtual target device, which maps the virtual SCSI server adapter to a file-backed virtual optical device, run the **mkvdev** command:

   ```
   mkvdev -fbo -vadapter VirtualSCSIServerAdapter
   ```

   where, *VirtualSCSIServerAdapter* is the name of the virtual SCSI server adapter. For example, vhost1.

   Note: No backing device is specified when you create virtual target devices for file-backed virtual optical devices because the drive is considered to contain no media. For information about loading media into a file-backed optical drive, see the **loadopt** command.

The optical device is available to the client logical partition either the next time it starts, or the next time the appropriate virtual SCSI client adapter is probed (on a Linux logical partition), or configured (on an AIX logical partition), or appears as an OPTXXX device (on an IBM i logical partition).

3. View the newly created virtual target device by running the **lsdev** command.

   For example, running `lsdev -virtual` returns results similar to the following:

   ```
   name     status     description
   vhost4   Available  Virtual SCSI Server Adapter
   vsa0     Available  LPAR Virtual Serial Adapter
   vtopt0   Available  Virtual Target Device - File-backed Optical
   ```

4. View the logical connection between the newly created devices by running the **lsmap** command.

   For example, running `lsmap -vadapter vhost1` returns results similar to the following:

   ```
   SVSA     Physloc                  Client PartitionID
   ----------------------------------------------------
   vhost1   U9117.570.10C8BCE-V6-C2  0x00000000

   VTD                vtopt0
   LUN                0x8200000000000000
   Backing device     Physloc
   ```

   The physical location is a combination of the slot number, in this case 2, and the logical partition ID. The virtual device can now be attached from the client logical partition.

What to do next

You can use the **loadopt** command to load file-backed virtual optical media into the file-backed virtual optical device.

If you later need to remove the virtual target device, you can do so by using the **rmvdev** command.

Related information

Creating a virtual disk for a VIOS logical partition using the HMC

Virtual I/O Server commands

Setting the reserve policy attributes of a device

In some configurations, you must consider the reservation policy of the device on the Virtual I/O Server (VIOS).

About this task

The following section explains the situations in which the reservation policy of the device on the VIOS is important for systems that are managed by the Hardware Management Console (HMC).

Situations where reservation policy of a device is important for HMC-managed systems

- To use a Multipath I/O (MPIO) configuration at the client, none of the virtual Small Computer Serial Interface (SCSI) devices on the VIOS can reserve the virtual SCSI device. Set the `reserve_policy` attribute of the device to `no_reserve`.
- For virtual SCSI devices used with Live Partition Mobility or the Suspend/Resume feature, the reserve attribute on the physical storage that is used by the mobile partition can be set as follows:
  - You can set the reserve policy attribute to `no_reserve`.
  - You can set the reserve policy attribute to `pr_shared` when the following products are at the following versions:
    - HMC Version 7 release 3.5.0, or later
    - VIOS Version 2.1.2.0, or later
    - The physical adapters support the SCSI-3 Persistent Reserves standard

  The reserve attribute must be the same on the source and destination VIOS partitions for successful partition mobility.

- For PowerVM Active Memory Sharing or Suspend/Resume features, the VIOS automatically sets the `reserve` attribute on the physical volume to `no reserve`. The VIOS performs this action when you add a paging space device to the shared memory pool.

Procedure

1. From a VIOS partition, list the disks (or paging space devices) to which the VIOS has access. Run the following command:

   ```
   lsdev -type disk
   ```

2. To determine the reserve policy of a disk, run the following command, where *hdiskX* is the name of the disk that you identified in step "1" on page 115. For example, hdisk5.

   ```
   lsdev -dev hdiskX -attr reserve_policy
   ```

   The results might look like the following output:

   ```
   ..
   reserve_policy  no_reserve                            Reserve Policy                  True
   ```

   Based on the information in the section Situations where reservation policy of a device is important for HMC-managed systems, you might need to change the reserve_policy so that you can use the disk in any of the described configurations.

3. To set the reserve_policy, run the **chdev** command.

   For example:

   ```
   chdev -dev hdiskX -attr reserve_policy=reservation
   ```

   where,

   - *hdiskX* is the name of the disk for which you want to set the reserve_policy attribute to `no_reserve`.
   - *reservation* is either `no_reserve` or `pr_shared`.

4. Repeat this procedure from the other VIOS partition.

   Requirements:

   a. Although the reserve_policy attribute is an attribute of the device, each VIOS saves the value of the attribute. You must set the reserve_policy attribute from both VIOS partitions so that both VIOS partitions recognize the reserve_policy of the device.

   b. For partition mobility, the reserve_policy on the destination VIOS partition must be the same as the reserve_policy on the source VIOS partition. For example, if the reserve_policy on the source VIOS partition is `pr_shared`, the reserve_policy on the destination VIOS partition must also be `pr_shared`.

   c. With the PR_exclusive mode on SCSI-3 reserve, you cannot migrate from one system to another system.

   d. The PR_key value for the VSCSI disks on the source system and the target system must be different.

Creating logical volume storage pools on a Virtual I/O Server

You can create a logical volume storage pool on a Virtual I/O Server by using the Hardware Management Console or the **mksp** and **mkbdsp** commands.

Before you begin

Before you start, ensure that the Virtual I/O Server is at Version 1.5 or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.

Tip: If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create logical volume storage pools on the Virtual I/O Server.

About this task

Logical volume storage pools are volume groups, which are collections of one or more physical volumes. The physical volumes that comprise a logical volume storage pool can be of varying sizes and types.

To create a logical volume storage pool, complete the following steps from the Virtual I/O Server command-line interface:

Procedure

1. Create a logical volume storage pool by running the **mksp** command:

   ```
   mksp -f dev_clients hdisk2 hdisk4
   ```

   In this example, the name of the storage pool is `dev_clients` and it contains `hdisk2` and `hdisk4`.
2. Define a logical volume, which will be visible as a disk to the client logical partition. The size of this logical volume will act as the size of disks that will be available to the client logical partition. Use the **mkbdsp** command to create an 11 GB logical volume called `dev_dbsrv` as follows:

   ```
   mkbdsp -sp dev_clients 11G -bd dev_dbsrv
   ```

   If you also want to create a virtual target device, which maps the virtual Small Computer Serial Interface (SCSI) server adapter to the logical volume, add `-vadapter vhost`*x* to the end of the command. For example:

   ```
   mkbdsp -sp dev_clients 11G -bd dev_dbsrv -vadapter vhost4
   ```

Related information

Creating storage pools on a Virtual I/O Server by using the HMC

Virtual I/O Server commands

Creating file storage pools on a Virtual I/O Server

You can create a file storage pool on a Virtual I/O Server by using the **mksp** and **mkbdsp** commands.

Before you begin

Before you start, ensure that the Virtual I/O Server is at Version 1.5 or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.

Tip: If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create file storage pools on the Virtual I/O Server.

About this task

File storage pools are created within a parent logical volume storage pool and contain a logical volume that contains a file system with files.

To create a file storage pool, complete the following steps from the Virtual I/O Server command-line interface:

Procedure

1. Create a file storage pool by running the **mksp** command:

   ```
   mksp -fb dev_fbclt -sp dev_clients -size 7g
   ```

   In this example, the name of the file storage pool is `dev_fbclt` and the parent storage pool is `dev_clients`.

2. Define a file, which will be visible as a disk to the client logical partition. The size of the file determines the size of the disk presented to the client logical partition. Use the **mkbdsp** command to create a 3 GB file called `dev_dbsrv` as follows:

   ```
   mkbdsp -sp dev_fbclt 3G -bd dev_dbsrv
   ```

   If you also want to create a virtual target device, which maps the virtual Small Computer Serial Interface (SCSI) server adapter to the file, add `-vadapter vhost`*x* to the end of the command. For example:

   ```
   mkbdsp -sp dev_fbclt 3G -bd dev_dbsrv -vadapter vhost4
   ```

Related information

Creating storage pools on a Virtual I/O Server by using the HMC

Virtual I/O Server commands

Creating the virtual media repository on a Virtual I/O Server

You can create the virtual media repository on a Virtual I/O Server with the **mkrep** command.

Before you begin

Before you start, ensure that the Virtual I/O Server is at Version 1.5 or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.

About this task

The virtual media repository provides a single container to store and manage file-backed virtual optical media files. Media that are stored in the repository can be loaded into file-backed virtual optical devices for exporting to client partitions.

Only one repository can be created within a Virtual I/O Server.

Tip: If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create a virtual media repository on the Virtual I/O Server.

Procedure

To create the virtual media repository from the Virtual I/O Server command-line interface, run the **mkrep** command:

```
mkrep -sp prod_store -size 6g
```

In this example, the name of the parent storage pool is `prod_store`.

Related information

Changing optical devices by using the Hardware Management Console

Virtual I/O Server commands

Creating volume groups and logical volumes on a Virtual I/O Server

You can create logical volumes and volume groups on a Virtual I/O Server by using the **mkvg** and **mklv** commands.

About this task

If you are using the HMC, Version 7 release 3.4.2 or later, you can use the HMC graphical interface to create volume groups and logical volumes on a Virtual I/O Server.

Otherwise, use the **mklv** command from the Virtual I/O Server command-line interface. To create the logical volume on a separate disk, you must first create a volume group and assign one or more disks by using the **mkvg** command.

Procedure

1. Create a volume group and assign a disk to this volume group by using the **mkvg** command. In this example, the name of the volume group is `rootvg_clients`

   ```
   mkvg -f -vg rootvg_clients hdisk2
   ```

2. Define a logical volume, which will be visible as a disk to the client logical partition. The size of this logical volume will act as the size of disks that will be available to the client logical partition. Use the **mklv** command to create a 2 GB logical volume as follows:

   ```
   mklv -lv rootvg_dbsrv rootvg_clients 2G
   ```

Related information

Changing a physical volume for a VIOS logical partition using the HMC

Changing a storage pool for a VIOS logical partition using the HMC

Configure the Virtual I/O Server to support SCSI-2 reserve functions

Understand the virtual Small Computer Serial Interface (SCSI) setup requirements to support applications by using SCSI reserve and release.

About this task

Virtual I/O Server Versions 1.3, and later provide support for applications that are enabled to use SCSI-2 reserve functions that are controlled by the client logical partition. Typically, SCSI reserve and release are used in clustered environments where contention for SCSI disk resources might require greater control. To ensure that Virtual I/O Server supports these environments, configure the Virtual I/O Server to support SCSI-2 reserve and release. If the applications you are using provide information about the policy to use for the SCSI-2 reserve functions on the client logical partition, follow those procedures for setting the reserve policy.

Complete the following tasks to configure the Virtual I/O Server to support SCSI-2 reserve environments:

Procedure

1. Configure the Virtual I/O Server reserve_policy for single_path, by using the following command:

   ```
   chdev -dev1 hdiskN -attr reserve_policy=single_path
   ```

   Note: Perform this task when the device is not in use. If you run this command while the device is open or in use, then you must use the **-perm** flag with this command. If you use the **-perm** flag, the changes do not take effect until the device is unconfigured and reconfigured.
2. Configure the client_reserve feature on the Virtual I/O Server.
   - If you are creating a virtual target device, use the following command:

     ```
     mkvdev -vdev hdiskN -vadapter vhostN -attr client_reserve=yes
     ```

     where, *hdiskN* is the virtual target device name and *vhostN* is the virtual SCSI server adapter name.
   - If the virtual target device has already been created, use the following command:

     ```
     chdev -dev vtscsiN -attr client_reserve=yes
     ```

     where, *vtscsiN* is the virtual device name.

Note: If the *client_reserve* attribute is set to *yes*, you cannot set the *mirrored* attribute to *true*. This is because the client_reserve and Peer-to-Peer Remote Copy (PPRC) features are mutually exclusive.

3. On the Virtual client, complete the following steps to configure the SCSI reserve and release support for the virtual disk backed by the physical disk that you configured in step 1. This is specific to an AIX client.

   a) Set the reserve policy on the Virtual client to single_path, by using the following command:

   ```
   chdev -a reserve_policy=single_path -1 hdiskN
   ```

   where, *hdiskN* is the virtual disk name

   Note: Perform this task when the device is not in use. If you run this command while the device is open or in use, then you must use the **-P** flag. In that case, the changes do not take effect until the device is unconfigured and reconfigured.

   b) Set the hcheck_cmd attribute so that the MPIO code uses the inquiry option. If the hcheck_cmd attribute is set to **test unit ready** and the backing device is reserved, then *test unit ready* fails and log an error on the client.

   ```
   chdev -a hcheck_cmd=inquiry -1 hdiskN
   ```

   where, *hdiskN* is the virtual disk name.

4. On the Virtual client, complete the following steps to configure the SCSI reserve and release support for the virtual disk backed by the physical disk that you configured in step 1. This is specific to an AIX client.

   a) Set the reserve policy on the Virtual client to single_path, by using the following command:

   ```
   chdev -a reserve_policy=single_path -1 hdiskN
   ```

   where, *hdiskN* is the virtual disk name

   Note: Perform this task when the device is not in use. If you run this command while the device is open or in use, then you must use the **-P** flag. In that case, the changes do not take effect until the device is unconfigured and reconfigured.

   b) Set the hcheck_cmd attribute so that the MPIO code uses the inquiry option. If the hcheck_cmd attribute is set to **test unit ready** and the backing device is reserved, then *test unit ready* fails and log an error on the client.

   ```
   chdev -a hcheck_cmd=inquiry -1 hdiskN
   ```

   where, *hdiskN* is the virtual disk name.

Configure the Virtual I/O Server to support exporting the PPRC secondary disk to client partitions

This topic describes how to export a Peer-to-Peer Remote Copy (PPRC) secondary device to a client partition. You can perform this task by creating a virtual target device with the PPRC secondary device as a backing device.

About this task

Virtual I/O Server (VIOS) Versions 2.2.0.0 and later, provide support for devices that are enabled to use the Peer-to-Peer Remote Copy (PPRC) feature. The PPRC feature can be used for real-time mirroring of disks. Typically, a PPRC pair consists of a primary virtual target device and a secondary virtual target device. The secondary virtual target device stores the backup data from the primary virtual target device. To enable exporting of the PPRC secondary virtual target device to a client partition, use the following command:

```
mkvdev -vdev hdiskN -vadapter vhostN -attr mirrored=true
```

Where,

- *hdiskN* is the secondary virtual target device name
- *vhostN* is the virtual Small Computer Serial Interface (SCSI) server adapter name

Identifying exportable disks

To export a physical volume as a virtual device, the physical volume must have an IEEE volume attribute, a unique identifier (UDID), or a physical identifier (PVID).

About this task

To identify exportable disks, complete the following steps:

Procedure

1. Determine whether a device has an IEEE volume attribute identifier by running the following command from the Virtual I/O Server command line:

   ```
   lsdev -dev hdiskX -attr
   ```

   Disks with an IEEE volume attribute identifier have a value in the `ieee_volname` field. Output similar to the following is displayed:

   ```
   ...
   cache_method   fast_write                       Write Caching method
      False
   ieee_volname   600A0B800012DD0D00000AB441ED6AC IEEE Unique volume name
      False
   lun_id         0x001a000000000000               Logical Unit Number
      False
   ...
   ```

   If the `ieee_volname` field does not appear, then the device does not have an IEEE volume attribute identifier.
2. If the device does not have an IEEE volume attribute identifier, then determine whether the device has a UDID by completing the following steps:
   a) Type `oem_setup_env`.
   b) Type `odmget -qattribute=unique_id CuAt`. The disks that have a UDID are listed.

      Output similar to the following is displayed:

      ```
      CuAt:
          name = "hdisk1"
          attribute = "unique_id"
          value = "2708ECVBZ1SC10IC35L146UCDY10-003IBXscsi"
          type = "R"
          generic = ""
          rep = "nl"
          nls_index = 79

      CuAt:
          name = "hdisk2"
          attribute = "unique_id"
          value = "210800038FB50AST373453LC03IBXscsi"
          type = "R"
          generic = ""
          rep = "nl"
          nls_index = 79
      ```

      Devices in the list that are accessible from other Virtual I/O Server partitions can be used in virtual Small Computer Serial Interface (SCSI) MPIO configurations.
   c) Type `exit`.
3. If the device does not have either an IEEE volume attribute identifier or a UDID, then determine whether the device has a PVID by running the following command:

```
lspv
```

The disks and their respective PVIDs are listed. Output similar to the following is displayed:

```
NAME        PVID                    VG       STATUS
hdisk0      00c5e10c1608fd80        rootvg   active
hdisk1      00c5e10cf7eb2195        rootvg   active
hdisk2      00c5e10c44df5673        None
hdisk3      00c5e10cf3ba6a9a        None
hdisk4      none                    None
```

4. If the device does not have either an IEEE volume attribute identifier, a UDID, or a PVID, then complete one of the following tasks to assign an identifier:

   a) Upgrade your vendor software and then repeat this entire procedure, Identifying exportable disks, from the beginning. The latest versions of some vendor software include support for identifying devices by using a UDID. Before upgrading, ensure that you preserve any virtual SCSI devices that you created when using the versions of the software that did not support identifying devices by using a UDID. For more information and upgrade instructions, see the documentation that is provided by your vendor software.

   b) If the upgraded vendor software does not produce a UDID or IEEE volume attribute identifier, then put a PVID on the physical volume by running the following command:

   ```
   chdev -dev hdiskX -attr pv=yes
   ```

Getting started with shared storage pools by using the VIOS command line interface

Learn about using the Virtual I/O Server (VIOS) command-line interface to manage shared storage pools.

On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, you can create a clustering configuration. The VIOS partition in a cluster is connected to the shared storage pool. VIOS partitions that are connected to the same shared storage pool must be part of the same cluster. Each cluster has a default storage pool. You can use the VIOS command-line interface to manage shared storage pools.

Notes:

- On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, a cluster consists of only one VIOS partition. VIOS Version 2.2.1.0 supports only one cluster in a VIOS partition.
- On VIOS Version 2.2.1.3, or later, a cluster consists of up to four networked VIOS partitions.
- On VIOS Version 2.2.2.0, or later, a cluster consists of up to 16 networked VIOS partitions. You can create a cluster with an Internet Protocol version 6 (IPv6) address that is configured on the VIOS logical partition.

In VIOS version 3.1, Shared Storage Pool (SSP) Management data is stored in the **PostgreSQL** database. All data files of the database are stored in the file system of the SSP cluster pool. If the VIOS node that manages the SSP database is unable to access the file system of the SSP cluster pool, while the **PostgreSQL** database process is performing an I/O operation, the **PostgreSQL** database aborts all operations and generates the core memory dump. The **PostgreSQL** database also generates the pool file system errors and stores them in the system error log file. The SSP database automatically recovers when the VIOS node that manages the SSP database regains access to the file system of the SSP cluster pool.

The following sections describe how you can create a configuration of a cluster with each cluster consisting of up to 16 VIOS partitions and several client partitions that use logical units, and how you can use the VIOS command-line interface.

To perform the shell command operations that are listed in the following sections on the VIOS, log in to the VIOS by using the **padmin** user ID.

Mirroring a shared storage pool

You can create, list, modify, or remove a failure group by using the command line interface on the Virtual I/O Server (VIOS) Version 2.2.3.0, or later.

Creating a failure group in a shared storage pool

You can create a mirrored copy of an existing shared storage pool.

- To create a failure group in the shared storage pool, run the **failgrp** command. Ensure that the size of the new failure group is more than or equal to the current pool size. In the following example, the *hdisk2* and *hdisk3* failure groups are used to create a mirrored copy of the shared storage pool:

  ```
  failgrp -create -clustername clusterA -sp poolA -fg FG1: hdisk2 hdisk3
  ```

- On VIOS Version 2.2.3.0, or later, you can create up to a maximum of two failure groups in the shared storage pool.

Listing failure groups in the shared storage pool

You can view the list of all the failure groups in the shared storage pool:

- To list all the failure groups in the shared storage pool, enter the following command:

  ```
  failgrp -list
  ```

- To change an existing failure group name in the shared storage pool, enter the following command:

  ```
  failgrp -modify -clustername clusterA -sp poolA -fg FG1 -attr name=newFG
  ```

- To check whether the failure group name is changed in the shared storage pool, enter the following command:

  ```
  failgrp -list -clustername clusterA -sp poolA
  ```

Removing an existing failure group

You can remove an existing failure group in the shared storage pool:

- To remove an existing failure group from the shared storage pool, enter the following command:

  ```
  failgrp -remove -clustername clusterA -sp poolA -fg Default
  ```

- To check whether the failure group name is removed from the shared storage pool, enter the following command:

  ```
  failgrp -list -clustername clusterA -sp poolA
  ```

Note: You cannot remove the failure group if only one failure group exists in a shared storage pool.

Configuring the system to create shared storage pools

Learn about configuring the system to create Virtual I/O Server (VIOS) shared storage pools.

Before you create shared storage pools, ensure that all logical partitions are preconfigured by using the Hardware Management Console (HMC) as described in this topic. The following are the supported number of characters for the names:

- Cluster: 63
- Storage pool: 127
- Failure group: 63
- Logical unit: 127

Configuring the VIOS logical partitions

Configure 16 VIOS logical partitions with the following characteristics:

- There must be at least one CPU and one physical CPU of entitlement.
- The logical partitions must be configured as a VIOS logical partitions.
- The logical partitions must consist of at least 4 GB of memory.
- The logical partitions must consist of at least one physical Fibre Channel adapter.
- The rootvg device for a VIOS logical partition cannot be included in storage pool provisioning.
- The associated rootvg device must be installed with VIOS Version 2.2.2.0, or later.
- The VIOS logical partition must be configured with sufficient number of virtual server Small Computer Serial Interface (SCSI) adapter connections required for the client logical partitions.
- The VIOS logical partitions in the cluster requires access to all the SAN-based physical volumes in the shared storage pool of the cluster.

The One VIOS logical partition must have a network connection either through an Integrated Virtual Ethernet adapter or through a physical adapter. On VIOS Version 2.2.2.0, clusters support virtual local area network (VLAN) tagging.

Note: In shared storage pools, the Shared Ethernet Adapter must be in threaded mode. For more information, see "Network attributes" on page 266.

Restriction: The VIOS logical partition must not be a mover service partition or a paging partition.

Restriction: You cannot use the logical units in a cluster as paging devices for PowerVM Active Memory Sharing or Suspend/Resume features.

Configuring client logical partitions

Configure client logical partitions with the following characteristics:

- The client logical partitions must be configured as AIX or Linux client systems.
- They must have at least 1 GB of minimum memory.
- The associated rootvg device must be installed with the appropriate AIX or Linux system software.
- Each client logical partition must be configured with enough virtual SCSI adapter connections to map to the virtual server SCSI adapter connections of the required VIOS logical partitions.

You can define more client logical partitions.

Storage provisioning

When a cluster is created, you must specify one physical volume for the repository physical volume and at least one physical volume for the storage pool physical volume. The storage pool physical volumes are used to provide storage to the actual data generated by the client partitions. The repository physical volume is used to perform cluster communication and store the cluster configuration. The maximum client storage capacity matches the total storage capacity of all storage pool physical volumes. The repository disk must have at least 1 GB of available storage space. The physical volumes in the storage pool must have at least 20 GB of available storage space in total.

Use any method that is available for the SAN vendor to create each physical volume with at least 20 GB of available storage space. Map the physical volume to the logical partition Fibre Channel adapter for each VIOS in the cluster. The physical volumes must be mapped only to the VIOS logical partitions connected to the shared storage pool.

Note: The Each of the VIOS logical partitions assigns *hdisk* names to all physical volumes available through the Fibre Channel ports, such as *hdisk0* and *hdisk1*. The VIOS logical partition might select different *hdisk* numbers for the same volumes to the other VIOS logical partition in the same cluster. For example, the *viosA1* VIOS logical partition can have *hdisk9* assigned to a specific SAN disk, whereas the *viosA2* VIOS logical partition can have the *hdisk3* name assigned to that same disk. For some tasks, the unique device ID (UDID) can be used to distinguish the volumes. Use the **chkdev** command to obtain the UDID for each disk.

Cluster communication mode

In VIOS 2.2.3.0 or later, by default, the shared storage pool cluster is created in a unicast address mode. In earlier VIOS versions, the cluster communication mode is created in a multicast address mode. When the cluster versions are upgraded to VIOS Version 2.2.3.0, the communication mode changes from multicast address mode to unicast address mode as part of rolling upgrade operation.

Related tasks

Migrating a cluster from IPv4 to IPv6
With the Virtual I/O Server (VIOS) Version 2.2.2.0, or later, you can migrate an existing cluster from Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6).

Related information

chkdev command

Failure group

Failure group refers to one or more physical disks that belong to one failure domain. When the system selects a mirrored physical partition layout, it considers the failure group as a single point of failure. For example, a failure group can represent all the disks that are the children of one particular adapter (adapterA versus adapterB), or all the disks that are present on one particular SAN (sanA versus sanB), or all the disks that are present in one particular geographic location (buildingA versus buildingB).

Shared storage pool mirroring

The data in the shared storage pool can be mirrored across multiple disks within a tier. In other words, it cannot be mirrored across tiers. The pool can withstand a physical disk failure by using the disk mirrors. During disk failures, SSP mirroring provides better reliability for the storage pool. Therefore, mirroring provides higher reliability and storage availability in the shared storage pool. The existing non-mirrored shared storage pool can be mirrored by providing a set of new disks that matches the capacity of the original failure group. All new disks belong to the new failure group.

If one or more disks or partitions of a mirrored pool fail, you receive alerts and notifications from the management console. When you receive alerts or notifications, you must replace the disk that failed with another functional disk. When the disk functions again or if the disk is replaced, the data is resynchronized automatically.

Networking considerations for shared storage pools

Learn about the networking considerations and restrictions for shared storage pools (SSP).

Networking considerations

The networking considerations for shared storage pools (SSP) follow:

- Uninterrupted network connectivity is required for SSP operations. The network interface that is used for the SSP configuration must be on a highly reliable network, which is not congested.
- Ensure that both the forward and the reverse lookup for the hostname that is used by the VIOS logical partition for clustering resolves to the same IP address.
- With the VIOS Version 2.2.2.0, or later, clusters support Internet Protocol version 6 (IPv6) addresses. Therefore, VIOS logical partitions in a cluster can have hostnames that resolve to an IPv6 address.
- To set up clusters on an IPv6 network, IPv6 stateless auto-configuration is recommended. You can have a VIOS logical partition that is configured with either IPv6 static configuration or IPv6 stateless auto-configuration. A VIOS logical partition that has both IPv6 static configuration and IPv6 stateless auto-configuration is not supported in VIOS Version 2.2.2.0.

- The hostname of each VIOS logical partition that belongs to the same cluster must resolve to the same IP address family, which is either Internet Protocol version 4 (IPv4) or IPv6 address.

Restrictions:

- In a cluster configuration, to change the hostname or the IP address of a VIOS logical partition, complete one of the following procedures depending on the number of VIOS logical partitions in the cluster:
  - If VIOS logical partitions exists in the cluster, remove the VIOS logical partition from the cluster and change the hostname or the IP address. You can later add the VIOS logical partition to the cluster again with the new hostname or the IP address.
  - If only one VIOS logical partition exists in the cluster, you must delete the cluster to change the hostname or the IP address. Before deleting the cluster, you must create a backup of SSP configuration by using the **viosbr** command. You can restore the cluster after the hostname or the IP address is updated.
- You must apply any hostname or the IP address changes to the /etc/netsvc.conf file of the VIOS logical partition before creating the cluster. This file is used to specify the order of name resolution for networking routines and commands. Later, if you want to edit the /etc/netsvc.conf file, complete the following procedure on each VIOS logical partition:
  1. To stop cluster services on the VIOS logical partition, type the following command:

     ```
     clstartstop -stop -n clustername -m vios_hostname
     ```

  2. Make the required changes in the /etc/netsvc.conf file. Do not change the IP address that resolves to the hostname that is being used for the cluster.
  3. To restart cluster services on the VIOS logical partition, type the following command:

     ```
     clstartstop -start -n clustername -m vios_hostname
     ```

  Maintain the same order of name resolution for all the VIOS logical partitions that belong to the same cluster. You must not make any changes to the /etc/netsvc.conf file when you are migrating a cluster from IPv4 to IPv6.

Multiple Transmission Control Protocol or Internet Protocol (TCP/IP) network support

In the VIOS versions earlier than VIOS Version 3.1.1.0, the shared storage pool (SSP) used only a single network interface or an IP for communication. Having a single network interface or an IP for communication might cause network failure and can be disruptive to the storage pool.

In the VIOS Version 3.1.1.0, or later, the shared storage pool improves the network resilience by supporting multiple TCP/IP network interfaces for LPAR client I/O specific communication. This communication is only used by the SSP for pool file system metadata protocol exchanges. Some of the VIOS daemon communication is also enhanced to use multiple network interfaces.

Multiple network interfaces are used in an active/passive mode. This means only one interface is used at a time without load balancing. In this case, one network interface is active and all the other network interfaces are in a standby mode. An active lease is maintained on all network interfaces for quick network interface switch-over. When the lease of an active network connection is at risk, the pool switches to another valid connection. The error log entries indicate the state of the network connection.

You can configure multiple TCP/IP network interfaces by using the **-addips** and **-rmips** options of the **cluster** command.

Best practices for using multiple TCP/IP networks:

- To achieve true redundancy of multiple TCP/IP networks, you must avoid using a single network interface for multiple network connections and configure separate isolated subnets.

- The network connection priority for multiple network connections is supported. In a multiple network interface environment, the primary network interface is utilized as much as possible. This means that if the primary network interface fails, failover to the secondary network interface occurs. Similarly, after the primary network is back and available, the communication automatically returns to the primary network interface. If the network interfaces have different speeds, the network interface with the highest speed must be defined as the primary network interface. For example, if the speed of the network interface is 10 gigabit and the speed of another network interface is 1 gigabit speed, the network interface with a 10 gigabit speed must be defined as the primary network interface. The IP address of this primary network interface resolves to the hostname that is used with a cluster node.
- Adding or removing the IP addresses when the node is online is not supported. You must stop the node to add or remove the network and then start the node again.
- Ensure that all the IP addresses of the cluster nodes are stored in the `/etc/hosts` file on all the nodes to avoid hostname query failure when the TCP/IP network or DNS is down. Failure in the hostname query might cause a node to take the shared storage pool offline on that node.

Limitations of multiple TCP/IP networks:

- Using the HMC to configure multiple IP addresses is not supported.
- You must stop and start the node for adding or removing the IP addresses. If you change the primary IP address or the hostname, remove the node from the cluster and then add it after the changes are complete.
- You can configure multiple network interfaces and create a backup of the network configuration by using the **viosbr** command. However, when you perform the complete cluster restore operation by using the backup file, the shared storage pool does not recognize any secondary interfaces. For the configured interfaces to be recognized, you must stop and start the node.
- The use of virtual IP address (VIPA) is not compatible while configuring multiple network interfaces by using the **cluster -addips** command. These are mutually exclusive techniques for network redundancy. The **cluster -addips** command cannot recognize a virtual IP address as it uses IP addresses from the physical network interfaces.

Disk communication support

In the VIOS Version 3.1.1.0 or later, you can configure disk communication for the shared storage pool LPAR client I/O specific communication. The shared storage pool keeps the disk connection active when all the TCP/IP networks are down. This allows you to manage a total network outage for a short period. The error log entry indicates when the node starts using disk communication and also when the network communication is resumed. When the TCP/IP network is back online, the shared storage pool automatically returns to communicate over the TCP/IP network.

A cluster is considered to be in a degraded mode when it is using disk communication:

- The primary goal of disk communication is to ensure that application I/O on client logical partitions (LPARs) do not time out.
- The VIOS CLI operations such as **cluster -status** might fail due to the network outage.
- Communication-intensive shared storage pool operations such as **PV remove** might also fail.

The communication disk is managed by Cluster Aware AIX (CAA) and it is separate from the repository disk. The size requirement for the disk is same as a repository disk. SSP supports only a single disk network for communication.

You can configure multiple TCP/IP network interfaces by using the **-addcompvs** and **-rmcompvs** options of the **cluster** command.

Best practices for disk communication:

- Provide a high-speed disk for disk communication depending on the I/O workload and the number of virtual I/O servers in the cluster.

- When an active TCP/IP network is not available, you cannot access the DNS. You must add the `/etc/hosts` entries for all the nodes to avoid a node getting expelled during the recovery operation and taking its pool offline.
- Disk communication is suited for low I/O rate applications such as **rootvg** or **middleware**. Disk communication can scale up to the limit of the storage performance.
- Reduce the application I/O operations during disk communication if the disk communication cannot handle the requests.
- During disk communication, you might need a larger error logging space for the `/var` and the `/home` directories when the networks are down. You need to monitor the `/var` and the `/home` directory space.

Limitations of disk communication:

- The database might not be accessible because it requires a TCP/IP network for connection.
- Configuration operations might fail because the database is not accessible.
- The **`cluster -status`** command might display that the shared storage pool is down because it does not use disk communication.
- 4K sector size disks are not supported for disk communication similar to repository disk.
- Using the HMC to configure disk communication is not supported.

Related information

cluster command

viosbr command

pv command

Adding flash acceleration to shared storage pools

Virtual I/O Servers (VIOS) with Shared Storage Pool (SSP) flash acceleration can improve performance by using Solid-State Drive (SSD) or flash storage caching on the Virtual I/O Server.

This feature enables each Virtual I/O Server to use a flash caching device for read-only caching. The flash caching devices can be:

- Devices that are attached to the server, such as a built-in SSD in the server.
- Devices that are directly attached to the server by using Serial Attached SCSI (SAS) controllers.
- Resources that are available in the storage area network (SAN).

The VIOS must be able to identify the device as a flash device for the device to be considered eligible to be used as a cache device. The VIOS uses the **MEDIUM ROTATION RATE** field of the **SCSI Block Device Characteristics VPD** page (SCSI INQUIRY page B1) to determine whether a device is a flash device. If a device does not support that page or displays a value other than *0001h Non-rotating medium* in the **MEDIUM ROTATION RATE** field, the device cannot be used as a cache device.

You can derive the maximum performance benefit by using locally attached flash caching devices.

SSP flash acceleration is based on caching on the Virtual I/O Servers, while Power flash caching or server-side caching is performed on the client logical partition. For more information on server-side caching, see Caching storage data or Integrated Server Based I/O Caching of SAN Based Data.

Both types of caching can be used independently. The performance characteristics of both of these types of caching are similar, on similar type of client logical partition workloads.

SSP flash acceleration performs read-only caching over the entire storage pool, including any storage tiers in the pool. Only the user data (data blocks) in the pool is cached, while the metadata is not cached. Instead, metadata access might be accelerated by using SSD storage on the SAN for the system tier.

Concepts and terms in SSP flash acceleration

You can cache the storage pool dynamically (enable or disable caching), while workloads are running on the client logical partitions. The workloads do not need to be brought down to an inactive state to enable caching. The terms that are used to explain the flash caching concept are described in the following table.

Term	Description
Cache device	A cache device is a Solid-State Drive (SSD) or a flash disk that is used for caching.
Cache pool	A cache pool is a group of cache devices that is used only for storage caching.
Enable caching	Start caching the storage pool.
Disable caching	Stop caching the storage pool.

When caching is enabled for the storage pool, caching starts on all Virtual I/O Servers in the cluster that have a defined cache pool. This process implicitly creates a logical cache device (known as a cache partition) derived from the local cache pool for each Virtual I/O Server. When the storage pool caching is enabled, all the read requests for the user data blocks of the storage pool are routed to the SSP caching software. If a specific user data block is found in the local Virtual I/O Server cache, the I/O request is processed from the cache device. If the requested block is not found in the cache, or if it is a write request, the I/O request is sent directly to the storage pool SAN devices.

When caching is disabled for the storage pool, the caching on all Virtual I/O Servers in the cluster stops. This process implicitly cleans up the logical cache device from the local cache pool on each server.

Architecture and components of SSP flash acceleration

The components of SSP flash acceleration include the VIOS, cache management and cache engine, and storage pool. These components are described in the following table.

Component	Description
VIOS	The administration and management of caching is performed from the VIOS command-line interface by using the **sspcache** command.
Storage pool (pool driver)	The storage pool is the caching target and the pool driver manages the cluster cache coherency.
Cache management and cache engine	Cache management provides the lower-level cache configuration commands, while the cache engine runs the local caching logic to determine what blocks are cached in the storage pool.

SSP flash acceleration performs distributed cache coherency between the Virtual I/O Servers in the following ways:

- The storage pool driver coordinates the distributed cache coherency across the cluster.
- The cache engine manages node level caching (promoting or demoting cache entries) and interacts with the storage pool driver to maintain cache coherency. This component uses the same local caching method as with the Power flash caching, or server-side caching.
- The cache engine is used for any storage pool I/O operations. This type of caching is known as *look-aside caching*.

The following figure explains the flow for various I/O operations when caching is enabled.

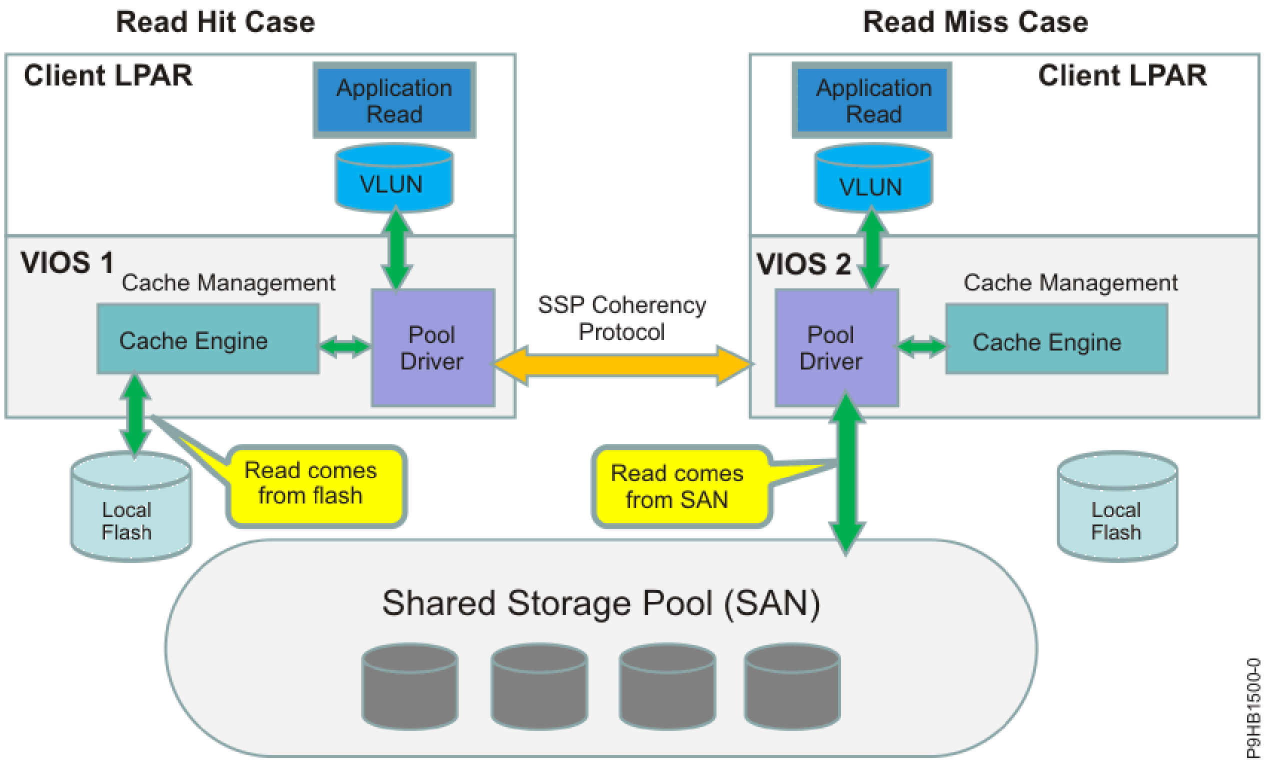

The details of the I/O operations that are shown in the figure, are explained in the following table.

I/O operation	Description
Cache Read Hit	• VIOS passes I/O read request from client logical partition to the storage pool driver. • Storage pool driver checks the cache engine and finds that the extent is cached in the local cache device. • The I/O request is entirely satisfied in the cache and passed back to the client logical partition.
Cache Read Miss	• VIOS passes I/O read request from client logical partition to the storage pool driver. • Storage pool driver checks the cache engine and finds that the extent is not cached in the local cache device. • The storage pool driver satisfies the request from the SAN and it is passed back to the client logical partition.
Write operation	• VIOS passes I/O write request from client logical partition to the storage pool driver. • The extent is invalidated on any node in the cluster that has the extent cached, to ensure cache coherency. • The storage pool driver performs the write request to the SAN.

Attributes of caching in SSP flash acceleration

The attributes of caching in SSP flash acceleration are:

Transparent to applications
Clustered applications can be used on the client logical partitions.

Independent of Client operating systems
Caching is supported on AIX, IBM i, and Linux operating systems.

Read only node-specific cache
Results of write operations are sent to the SAN after cache invalidation occurs.

Concurrent and coherent shared data access
Supports concurrent shared data access with full coherency across the SSP landscape.

Independent of types of storage
No dependency on the type of flash storage for caching and SAN storage for SSP.

Advantages of SSP flash acceleration

Some of the benefits of SSP flash acceleration include:

- Improvement in latency and throughput with certain workloads such as analytical and transactional workloads, and online transaction processing.
- Transparent acceleration, such that client logical partitions are unaware of caching on Virtual I/O Servers.
- Better virtual machine (VM) density, without performance impacts.
- Allows more efficient utilization and scaling of SAN infrastructure. The SAN offloading of read requests can increase write throughput on congested SANs.
- Benefits from sharing blocks across VMs based on cloned virtual Logical Units (LUs), when common blocks are already cached.
- Compatibility with Live Partition Mobility (LPM).

Limitations of caching in SSP flash acceleration

Some limitations of caching in SSP flash acceleration are:

- The SSP caching software is configured as a read-only cache, which means that only read requests are processed from the flash Solid-State Drive (SSD). All write requests are processed by the storage pool only and go directly to the SAN.
- Data that is written to the storage pool is not populated in the cache automatically. If the write operation is performed on a block that is in the cache, the existing data in the cache is marked as invalid. The same block reappears in the cache, based on how frequently and how recently the block is accessed.
- Cache devices cannot be shared between Virtual I/O Servers.
- Performance benefits depend on the size of the application working set and the type and size of the SAN disk controller cache. Typically, the collective working set must be larger than the SAN disk controller cache to realize significant performance benefits.

Configuration of caching in SSP flash acceleration

You must complete the following steps from the VIOS command-line interface to enable caching:

1. Create a cache pool on each VIOS in the cluster, by using the **`cache_mgt`** command.
2. Enable caching of the storage pool on the SSP cluster from a single VIOS node by using the **sspcache** command.

Creation of the cache pool on each VIOS is a one-time step. The syntax for this command is:

```
cache_mgt pool create –d <devName>[,<devName>,…] -p <poolName>
```

For example, to create a 1024 MB cache on each VIOS in the cluster and then to enable caching on the storage pool, complete the following steps:

1. To create a 1024 MB cache, enter the following command:

```
cache_mgt pool create –d /dev/hdisk11 –p cmpool0
```

This command must be run on all Virtual I/O Servers in the cluster.

2. To enable caching of the storage pool on the SSP cluster from a single VIOS node, enter the following command:

   ```
   sspcache -enable -sp -size 1024
   ```

   This command must be run on a single VIOS in the cluster.

Management of caching in SSP flash acceleration

After caching is configured, the caching requirements might change over time. You might need to add new workloads that need to be cached. To fulfill the changing requirements, the cache pool can be extended, by adding extra cache devices, if necessary. Thus, you can increase the cache size.

You can use the following examples to manage the caching configuration.

1. To add a cache device to the cache pool, enter the following command on each VIOS in the cluster:

   ```
   > cache_mgt pool extend -p cmpool0 -d hdisk12 –f
   ```

2. To extend the cache size to 2048 MB, enter the following command on one node:

   ```
   > sspcache –resize –sp –size 2048
   ```

Related information

sspcache command

Managing storage tiers

You can use the command-line interface on the Virtual I/O Server (VIOS) to manage a storage tier. You can also use the Hardware Management Console (HMC) version 8.4.0, or later to manage storage tiers.

Creating a storage tier

You can create a storage tier by using the VIOS command-line interface. The system tier is created when you create a cluster. This procedure focuses on creating a user tier.

Before you begin

Consider the following prerequisites for creating a user tier:

- A value must be entered for a failure group. If no value is specified for the failure group, a default failure group is created and named with the default failure group name *default*.
- A list of cluster capable physical volumes (PVs) must be available.
- The tier name must be limited to 63 characters, can be alphanumeric, and can contain '_' (underscore), '-" (hyphen), or '.' (dot) characters. The tier name must always be unique within a pool.

About this task

You can create a user tier with a specified name, containing specific physical volumes (PVs). PVs can be specified as a space-delimited list of PVs in the command, or as a file containing a space-delimited list of PVs. You can add user tiers to a pool by using the VIOS command-line interface (CLI). You can add only one tier at a time.

To create a tier with one VIOS logical partition, complete the following steps:

Procedure

1. Identify the PVs that you want to add to the new user tier.

2. To create a tier, run the following command: `tier -create -tier` *`tier1`*`:` *`hdisk1 hdisk2`*.

Examples:

1. To create a user tier that specifies the physical volumes to be used in a space-delimited list file, rather than in an inline list of physical volumes, enter the following command:

```
tier -create -file -tier tier1: /tmp/pvfile
```

2. To create a cluster with a restricted system tier and a separate user tier, enter the following command:

```
cluster –create –clustername cname –repopvs hdisk4
–sp pname –systier hdisk5 –usrtier userTier1:hdisk6
```

The system tier is automatically marked as *Restricted* while the user tier is automatically marked as *default*.

Related information

tier command

Setting the storage tier type

A system tier must be identified as either a restricted system tier or a co-mingled tier. You can set the tier type by using the Virtual I/O Server (VIOS) command-line interface.

About this task

When you create a cluster without any command parameters, a co-mingled tier (type *comingled*) is created by default. Co-mingled tier contains both the metadata and the user data. If you want to separate the metadata from the user data, you can change the type of the tier to *system*.

To change the type of a tier to *restricted* system tier, enter the following command:

```
tier -modify -tier SYSTEM -attr type=system
```

The **tier -modify** command with the *-attr* value *system* can be used only for system tiers and cannot be used for user tiers. In this example, the system tier with the name *SYSTEM* is now set as a *restricted* system tier.

Setting the default storage tier

A default storage tier must be identified within a storage pool. The default tier is created first. You can change the default storage tier by using the Virtual I/O Server (VIOS) command-line interface.

About this task

The first user data tier that is created during the creation of the cluster, is the default (or provisioning) tier. This is the default tier, only for the placement of user data for virtual disks if a tier name is not specified. The default tier for user data can be changed, if the chosen default is not suitable.

To set a storage tier as the default storage tier, complete the following steps:

```
Enter the following command: tier -modify -tier tier1 -attr default=yes.
```

The storage tier with the name *tier1* is now set as the default tier. Because you can have only one default tier, the setting for the previous default tier is automatically set to`default=`*no*.

Listing the storage tiers

You can list the existing storage tiers in a storage pool by using the Virtual I/O Server (VIOS) command-line interface.

About this task

If you need to add a logical unit (LU) or physical volume (PV) to an existing storage tier within a storage pool, list the names and details to determine which tier has the available space, or to determine which tier you want to add the LU or PV.

To list the tiers, complete the following step:

Procedure

1. Enter the following command: `tier -list`.

   The following information about the storage tiers is provided within that storage pool:

   POOL_NAME
   The name of the storage pool.

   TIER_NAME
   The name of the tier that the information applies to.

   SIZE(MB)
   The size of the specified tier.

   FREE_SPACE(MB)
   The amount of free space that is available in the specified tier.

   MIRROR_STATE
   The current state of the mirroring activity on the specified tier, if applicable.
2. You can also list additional details for each tier by entering the following command: `tier -list -verbose`.

   In addition to the information that is provided by the **`tier -list`** command, the following information is also displayed:

   TIER_TYPE
   Whether the tier is a co-mingled tier, a user tier, or a restricted tier.

   TIER_DEFAULT
   Whether the tier is set as the default tier.

   OVERCOMMIT_SIZE
   The amount of space that can be used when the size of the tier is exceeded.

   TOTAL_LUS
   The number of LUs that are currently assigned to the tier.

   TOTAL_LU_SIZE
   The size in MB of all of the LUs that are assigned to that tier.

   FG_COUNT
   The number of failure groups that are assigned to that tier.

   ERASURE_CODE
   The identification of any mirrored tiers, if applicable.

Renaming a storage tier

You can rename a storage tier by using the Virtual I/O Server (VIOS) command-line interface.

About this task

All storage tiers must have a name for identification. Only the automatically created system tier is given a default name, which is SYSTEM. When you rename a shared storage pool tier, ensure that the new name has a maximum of 63 characters. The supported characters for the name are alphanumeric characters, - (dash), _ (underscore), or . (dot). To rename an existing storage tier, complete the following steps:

Procedure

1. Enter the following command: `tier -modify` *`oldTierName`* `-attr name=`*`newTierName`*
2. Enter the following command to verify that the tier was renamed: `tier -list`.

   The storage tier name now appears as *newTierName*.

Removing a storage tier

You can remove a storage tier by using the Virtual I/O Server (VIOS) command-line interface. You can remove only a system tier by removing the cluster.

Before you begin

Ensure that you understand and comply with the following restrictions before removing a tier:

- The tier must be empty. This means that any operation to move LUs out of the tier must have completed successfully. No LUs must be assigned to the tier and all LU blocks must be freed or migrated to other tiers successfully.
- You can remove only user tiers.
- You cannot remove the default storage tier. To remove a tier that is identified as the default tier, you must change the default tier to a different tier by using the tier command.
- You cannot remove the system tier. The only way to remove the system tier is to delete the cluster by using the `cluster -remove` command.

Procedure

To remove a tier, enter the following command: `tier -remove -tier tier1`.

This removes the *tier1* tier.

Managing a cluster by using the VIOS command line

You can use the command-line interface on the Virtual I/O Server (VIOS) to manage a cluster and the VIOS logical partitions.

Note: To add or remove devices in a cluster, you must use the Fully Qualified Domain Name (FQDN) of the device.

Creating a cluster with a single VIOS logical partition

You can create a cluster with a single VIOS logical partition by using the VIOS command-line interface.

Before you begin

Before you start, ensure that the following requirements are satisfied:

1. Log in to the `viosA1` VIOS logical partition by using the **padmin** user ID, which provides a restricted Korn shell environment.
2. Locate the physical volumes to be used for the `clusterA` cluster. For example, entering the `lspv -free` command returns results similar to the following:

```
NAME      PVID                    SIZE (megabytes)
-------------------------------------------------------
hdisk0    none                     17408
hdisk2    000d44516400a5c2         20480
hdisk3    000d4451605a0d99         20482
hdisk4    none                     10250
hdisk5    none                     20485
hdisk6    none                     20490
hdisk7    none                     20495
```

```
hdisk8    none                     20500
hdisk9    none                     20505
```

The **lspv** command displays a list of physical volumes along with the ID. The physical volume ID indicates that the device might be in use. The system administrator must ensure that the physical volume is not in use before you add it to the cluster repository or shared storage pool. For example, you can select the hdisk9 physical volume for the repository, and hdisk5 and hdisk7 physical volumes for the storage pool.

About this task

To create a cluster with one VIOS logical partition, complete the following steps:

Procedure

1. Run the **cluster** command to create the cluster.

   In the following example, the storage pool for the clusterA cluster is named poolA.

   ```
   cluster -create -clustername clusterA -repopvs hdisk9 -spname poolA -sppvs  hdisk5  hdisk7
   -hostname
   viosA1_HostName
   ```

2. After you create the cluster, run the **lspv** command to display the list of all the physical volumes visible to the logical partition.

   For example, entering the lspv command returns results similar to the following:

   ```
   NAME          PVID                    VG                STATUS
   ---------------------------------------------------------------------
   hdisk0        none                     None
   hdisk1        000d4451b445ccc7         rootvg            active
   hdisk2        000d44516400a5c2         20480
   hdisk3        000d4451605a0d99         10250
   hdisk4        none                     20485
   hdisk5        none                     20490
   hdisk6        none                     20495
   hdisk7        none                     20500
   hdisk8        none                     20505
   hdisk9        none                     caavg_private     active
   ```

   Note: The disk for the repository has a volume group name caavg_private. Volume group commands such as **exportvg** and **lsvg** must not be run on the repository disk.

3. To display a list of physical volumes for which the usage cannot be determined, run the **lspv** command.

   For example, entering the lspv -free command returns results similar to the following:

   ```
   NAME      PVID                    SIZE (megabytes)
   -------------------------------------------------------
   hdisk0    none                     17408
   hdisk2    000d44516400a5c2         20480
   hdisk3    000d4451605a0d99         20482
   hdisk4    none                     10250
   hdisk6    none                     20490
   hdisk8    none                     20500
   ```

4. To display the physical volumes in the storage pool, run the **lspv** command.

   For example, entering the lspv -clustername clusterA -sp poolA command returns results similar to the following:

   ```
   PV NAME      SIZE(MB)             PVUDID
   -------------------------------------------------------
   hdisk5       20480         200B75CXHW1026D07210790003IBMfcp
   hdisk7       20495         200B75CXHW1020207210790003IBMfcp
   ```

5. To display cluster information, run the **cluster** command.

For example, entering the `cluster -status -clustername clusterA` command returns results similar to the following:

```
Cluster Name         State
clusterA             OK

    Node Name        MTM           Partition Num  State  Pool State
    viosA1           9117-MMA0206AB272        15  OK     OK
```

What to do next

To list cluster configuration information, use the **lscluster** command. For example, entering the **lscluster -m** command returns results similar to the following:

```
Calling node query for all nodes
Node query number of nodes examined: 1

Node name: viosA1
Cluster shorthand id for node: 1
uuid for node: ff8dd204-2de1-11e0-beef-00145eb8a94c
State of node:  UP  NODE_LOCAL
Smoothed rtt to node: 0
Mean Deviation in network rtt to node: 0
Number of zones this node is a member in: 0
Number of clusters node is a member in: 1
CLUSTER NAME       TYPE  SHID   UUID
clusterA           local        a3fe209a-4959-11e0-809c-00145eb8a94c
Number of points_of_contact for node: 0
Point-of-contact interface & contact state
n/a
```

For more information, see lscluster command.

Related information

cluster command

lspv command

Replacing a repository disk

On the Virtual I/O Server (VIOS) Version 2.2.2.0, you can replace a repository disk by using the VIOS command-line interface.

About this task

You can replace the repository disk that is used to store cluster configuration information thus, increasing cluster resiliency. The replace operation works on a functional or failed repository disk. When the repository disk fails, the cluster remains operational. While the repository disk is in a failed state, all requests for cluster configuration fail. After you replace the failed disk, the cluster will be fully functional. As part of the replace operation, the cluster configuration information is stored on the new repository disk. The following are the requirements that must be met:

- The new repository disk must be at least the same size as the original disk.
- The VIOS logical partitions in the cluster must be at Version 2.2.2.0, or later.

Procedure

To replace a repository disk, run the **chrepos** command.

In the following example, the `hdisk1` repository disk is replaced with the `hdisk5` repository disk.

```
chrepos -n -r +hdisk5 -hdisk1
```

Adding a VIOS logical partition to a cluster

You can add a VIOS logical partition to a cluster by using the VIOS command-line interface.

About this task

To add a VIOS logical partition to a cluster:

Procedure

1. Run the **cluster** command to add a VIOS logical partition to a cluster. The fully qualified network host name for the VIOS logical partition must be specified. For example,

   ```
   cluster    -addnode -clustername clusterA     -hostname viosA2
   ```

   In this example, the viosA2 VIOS logical partition is added to the clusterA cluster.
2. To display the VIOS logical partitions in the cluster, use the **cluster** command.

   For example,

   ```
   cluster -status -clustername clusterA
   ```

3. Log in to the VIOS logical partition by using the **padmin** user ID to confirm the cluster characteristics as seen by the VIOS logical partition by entering the following commands:

   ```
   cluster -status -clustername clusterA

   lssp -clustername clusterA
   lssp -clustername clusterA -sp poolA -bd
   lspv -clustername clusterA -sp poolA
   ```

4. You can map the existing logical units to the virtual server adapters of the VIOS logical partitions.

   In this example, the logical units added to the viosA1 VIOS logical partition must be visible. However, these logical units are not yet mapped to the virtual server adapters that are provided by the viosA2 VIOS logical partition. To map existing logical units to the virtual server adapters of the viosA2 VIOS logical partition (while logged in to the viosA2 VIOS logical partition) and to list the mappings, enter the following commands:

   ```
   mkbdsp -clustername clusterA -sp poolA -bd luA1 -vadapter vhost0
   ```

   ```
   mkbdsp -clustername clusterA -sp poolA -bd luA2    -vadapter vhost1
   ```

   ```
   lsmap -clustername clusterA -all
   ```

   The client systems can now be reconfigured to accommodate the new mappings.

Related information

cluster command

lsmap command

lspv command

lssp command

mkbdsp command

Removing a VIOS logical partition from a cluster

You can remove a VIOS logical partition from a cluster by using the VIOS command-line interface.

About this task

After you add a logical partition to a cluster and enabling the client mapping to the same logical unit, you can remove the VIOS logical partition from the cluster. To remove a VIOS logical partition from a cluster:

Procedure

1. Run the **cluster** command to remove a VIOS logical partition from a cluster. Specify the fully qualified network host name for the VIOS logical partition.

   For example,

   ```
   cluster    -rmnode -clustername clusterA -hostname viosA1
   ```

   Note: You cannot run this command on the VIOS logical partition that is being removed.
2. To verify the removal of the node and the retention of objects that are still logged in to other partitions, run the **cluster** and **lssp** commands. For example,

   ```
   cluster -status -clustername clusterA

   lssp -clustername clusterA -sp poolA -bd
   lssp -clustername clusterA
   lspv -clustername clusterA -sp poolA
   ```

   Note: If the VIOS logical partition is mapped to a logical unit in the storage pool of the cluster, removing that VIOS logical partition from a cluster fails. To remove the logical partition, unmap the logical unit.

Related tasks

Unmapping a logical unit
You can unmap a logical unit by using the Virtual I/O Server (VIOS) command-line interface.

Related information

cluster command

lspv command

lssp command

Deleting a cluster

You can delete a cluster by using the Virtual I/O Server (VIOS) command-line interface.

About this task

Notes:

- You cannot restore a cluster if you delete the cluster. You cannot restore a VIOS logical partition in a cluster if the VIOS logical partition is removed from the cluster.
- Deleting a cluster fails if the VIOS logical partition has any mappings to logical units in the shared storage pool or if there are any logical units within the shared storage pool. Before you perform the delete operation, remove all logical partition mappings and logical units.

To delete a cluster, including the physical volumes provisioned to its storage pool, complete the following steps:

Procedure

1. Run the **cluster** command. For example, enter `cluster -delete -clustername clusterA` to delete the *clusterA* cluster.
2. To verify that the physical volumes are released to the free state, run the **lspv** command.

   For example, when you enter `lspv -free`, all the physical volumes must be displayed in the free physical volume list.

Related concepts

Removing logical units

You can remove logical units from the shared storage pool by using the Virtual I/O Server (VIOS) command-line interface.

Related tasks

Unmapping a logical unit
You can unmap a logical unit by using the Virtual I/O Server (VIOS) command-line interface.

Related information

cluster command

lspv command

Migrating a cluster from IPv4 to IPv6

With the Virtual I/O Server (VIOS) Version 2.2.2.0, or later, you can migrate an existing cluster from Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6).

Before you begin

Notes:

- You must not change the IP address of a VIOS logical partition in a cluster that resolves to the host name dynamically.
- You can migrate an existing cluster that is using IPv4 addresses to a cluster that is using IPv6 addresses only after each of the VIOS logical partitions are updated to VIOS Version 2.2.2.0, or later.

About this task

To migrate a cluster from IPv4 to IPv6:

Procedure

1. On the VIOS command line, type **mktcpip** to add an IPv6 address to each of the VIOS logical partitions that are in the IPv4 cluster. For more information about the commands that are used to configure an IPv6 address on the VIOS logical partition, see "Configuring IPv6 on the Virtual I/O Server" on page 204.

   Note: Do not remove the IPv4 addresses that the host name of each VIOS logical partition are resolving to until after you complete step 2 for all VIOS logical partitions.
2. Complete the following steps on each VIOS logical partition in the cluster:
   a) Stop cluster services on the VIOS logical partition by running the following command:

   ```
   clstartstop -stop -n clustername -m node_hostname
   ```

   b) Make the required changes in the network configuration, Neighbor Discovery Protocol (NDP) daemon router, or Domain Name System (DNS) information so that the IPv6 address of the VIOS logical partition resolves to the same host name that earlier resolved to the IPv4 address. Ensure that both the forward and reverse lookup for the same host name resolves to the required IPv6 address.

   c) On the VIOS command line, type the following command to restart cluster services on the VIOS logical partition:

   ```
   clstartstop -start -n clustername -m node_hostname
   ```

   d) Repeat steps 2a - 2c for each VIOS logical partition that belongs to the cluster.
3. From the VIOS command line, type **rmtcpip** to remove the IPv4 address from each VIOS logical partition.

Managing storage pools by using the VIOS command line

You can use the command-line interface on the Virtual I/O Server (VIOS) to manage shared storage pools.

Adding storage space to the storage pool

When more storage space is required in a storage pool, you can add one or more physical volumes in the storage pool by using the Virtual I/O Server (VIOS) command-line interface.

Adding physical volumes to the storage pool

You can add physical volumes to the storage pool by using the Virtual I/O Server (VIOS) command-line interface.

Before you begin

Prerequisites

Before you start, ensure that there are physical volumes capable of being added to the storage pool. To display a list of physical volumes for which the usage cannot be determined, enter the `lspv -free` or `lspv -capable` commands immediately before you change the storage provisioning. Another VIOS logical partition might have taken a physical volume. For example, entering the `lspv -free` command returns results similar to the following:

```
NAME      PVID                    SIZE (megabytes)
-----------------------------------------------------
hdisk0    none                     17408
hdisk3    000d4451605a0d99         20482
hdisk4    none                     10250
hdisk6    none                     20490
hdisk8    none                     20500
```

List the physical volumes that are capable of being included in the storage pool. For example, entering the `lspv -clustername clusterA -capable` command returns results similar to the following:

```
PV NAME                  SIZE (MB)      PVUDID
--------------------------------------------------------------------------
hdisk0                   17408          200B75CXHW1025F07210790003IBMfcp
hdisk3                   20482          200B75CXHW1031007210790003IBMfcp
hdisk4                   10250          200B75CXHW1031107210790003IBMfcp
hdisk6                   20490          200B75CXHW1031307210790003IBMfcp
hdisk8                   20500          200B75CXHW1031A07210790003IBMfcp
```

To determine if a physical volume is in use, run the **prepdev** command. If the physical volume is in use as a cluster repository disk or as a storage pool disk, you receive an error message. For example, by entering `prepdev -dev hdisk5`, you can determine if the *hdisk5* physical volume is in use. Output similar to the following is displayed:

```
WARNING!
The VIOS has detected that this physical volume is currently in use. Data will be
lost and cannot be undone when destructive actions are taken. These actions should
only be done after confirming that the current physical volume usage and data are
no longer needed.
The VIOS could not determine the current usage of this device.
```

If the physical volume is in use as a cluster repository disk or as a storage pool disk, you can use the **cleandisk** command to make the physical volume available.

Note: Ensure that the physical volume is no longer required, because running the **cleandisk** command results in loss of data on the physical volume.

- To remove a cluster repository disk signature from the *hdisk4* physical volume, enter the following command:

  ```
  cleandisk -r hdisk4
  ```

- To remove a storage pool disk signature from the *hdisk4* physical volume, enter the following command:

  ```
  cleandisk -s hdisk4
  ```

About this task

To add one or more physical volumes to a storage pool, complete the following steps:

Procedure

1. Add physical volumes to the storage pool by using the **chsp** command. For example,

   ```
   chsp -add -clustername clusterA -sp poolA hdisk4 hdisk8
   ```

   In this example, the *hdisk4* and *hdisk8* physical volumes are added to the storage pool.
2. To display the list of physical volumes in the storage pool, use the **lspv** command. For example, entering the `lspv -clustername clusterA -sp poolA` command returns results similar to the following:

   ```
   PV NAME                  SIZE (MB)     PVUDID
   --------------------------------------------------------------------------
   hdisk4                   20485         200B75CXHW1031207210790003IBMfcp
   hdisk5                   20495         200B75CXHW1031907210790003IBMfcp
   hdisk6                   10250         200B75CXHW1031107210790003IBMfcp
   hdisk8                   20500         200B75CXHW1031A07210790003IBMfcp
   ```

3. To display the list of the remaining free physical volumes that can be included in the cluster, use the **lspv** command.

   For example, entering the `lspv -clustername clusterA -capable` command returns results similar to the following:

   ```
   PV NAME                  SIZE (MB)     PVUDID
   --------------------------------------------------------------------------
   hdisk0                   17408         200B75CXHW1025F07210790003IBMfcp
   hdisk3                   20482         200B75CXHW1031007210790003IBMfcp
   hdisk6                   20490         200B75CXHW1031307210790003IBMfcp
   hdisk9                   20505         200B75CXHW1031A07210790003IBMfcp
   ```

4. To display the information about the shared storage pool, such as pool size, available free space, and how overcommitted the shared storage pool is, use the **lssp** command.

   For example, entering the `lssp -clustername ClusterA` command returns results similar to the following:

   ```
   POOL_NAME:       poolA
   POOL_SIZE:       71730
   FREE_SPACE:      4096
   TOTAL_LU_SIZE:   80480
   OVERCOMMIT_SIZE: 8750
   TOTAL_LUS:       20
   POOL_TYPE:       CLPOOL
   POOL_ID:         FFFFFFFFAC10800E000000004F43B5DA
   ```

5. To remove a physical volume (PV) from a shared storage pool (SSP), use the **pv** command. For more information, see the pv Command.

Related information

chsp command

cleandisk command

lspv command

prepdev command

Replacing physical volumes in the storage pool

You can replace physical volumes in the storage pool by using the command-line interface on VIOS Version 2.2.1.3, or later.

About this task

When more storage space is needed in a storage pool, you can also add or replace existing physical volumes in a storage pool. If you are replacing the existing physical volume with a physical volume that has a larger capacity, the capacity of the shared storage pool increases.

Restrictions:

- You can replace physical volumes only in one cluster at a time.
- Do not use this task to increase only the capacity of the shared storage pool.

To remove and replace physical volumes in the storage pool, complete the following steps:

Procedure

1. Remove and replace a physical volume that is in a storage pool by running the **chsp** command.

   For example,

   ```
   chsp -replace -clustername clusterA -sp poolA -oldpv hdisk4 -newpv hdisk9
   ```

   In this example, the `hdisk4` physical volume is replaced by the `hdisk9` physical volume in the storage pool. The replaced disk is returned to the free physical volume list.

   Note: If the size of physical volume that is being replaced is large, the replace operation might take a longer time to complete.

2. To see the new set of physical volumes in the storage pool, run the **`lspv`** command.

   For example, entering the `lspv -clustername clusterA -sp poolA` command returns results similar to the following:

   ```
   PV NAME                  SIZE (MB)       PVUDID
   ---------------------------------------------------------------------------
   hdisk0                   20485           200B75CXHW1031207210790003IBMfcp
   hdisk1                   20495           200B75CXHW1031907210790003IBMfcp
   hdisk8                   20500           200B75CXHW1031A07210790003IBMfcp
   hdisk9                   20505           200B75CXHW1031A07210790003IBMfcp
   ```

Related information

chsp command

lspv command

Changing the storage threshold

You can change the threshold limit of the storage usage by using the Virtual I/O Server (VIOS) command-line interface.

The shared storage pool space is used to store virtual client partition user data. You must view threshold alerts to verify if the free space decreases to a value that is lesser than the acceptable value.

Important: Free space must not reduce to a value that is lesser than 5% of the total space. If this reduction occurs, I/O operations on the virtual client partition might fail. To avoid this failure, you must add physical volumes to the pool or delete data from the pool to create free space.

The threshold limit for alert generation is a percentage value. If the actual storage usage transitions to a value that is either greater or lesser than the threshold limit, an alert is raised and an entry is made into the VIOS error log in the VIOS logical partition that is a Primary Notification Node (PNN). If a PNN does not exist, the error log is created on the Database Node (DBN). To determine whether the VIOS logical partition is a PNN or the DBN, run the **`lssrc -ls vio_daemon`** command. The system error log is used to track the threshold condition. These conditions are propagated to the Hardware Management Console (HMC) if they are connected to the VIOS partition. The threshold limit can be changed to a value from 1%

- 99%, with the number representing the amount of free space. The default threshold monitoring is set to alert when the free space decreases to a value that is lesser than 35% of the total capacity.

For example, if the threshold limit is 20% and the amount of free space decreases to a value that is lesser than 20%, an alert is raised with an indication that the threshold limit was exceeded. After you add storage space, by adding storage capacity to the storage pool, and the amount of free space exceeds 20%, another alert is raised with the indication that the threshold is no longer exceeded. An optimum threshold limit depends on the administrative capability to respond to alerts and on how quickly storage is used.

The following list describes how to change the threshold limit, and remove and view threshold alerts:

- To change the threshold limit, run the **`alert`** command. In the following example, the threshold limit is changed to 10%. Thus, an *exceeded* alert is raised when the free space decreases to a value that is lesser than 10% of the physical storage pool capacity.

  ```
  alert -set -clustername clusterA -spname poolA -type threshold -value 10
  ```

  Note: You can check threshold alerts in the VIOS system error log.

- To remove the threshold alert from the storage pool, enter the `alert -unset` command.

  ```
  alert -unset -clustername clusterA -spname poolA -type threshold
  ```

  Note: If you disable the threshold alert notification feature, a threshold alert will not be raised before the free space in a storage pool decreases to a value that is lesser than the acceptable value. Threshold alerts are important when you use thin-provisioned logical units in shared storage pool.

- To view the threshold alert on the storage pool, enter the `alert -list` command.

  ```
  alert -list -clustername clusterA -spname poolA -type threshold
  ```

- To list the error log, enter the `errlog –ls | more` command. You can look for log entries containing the following information:
  - Information messages
  - **VIO_ALERT_EVENT** label
  - *Threshold Exceeded* alert

The following list describes how to change the overcommit limit of the storage pool, view, and remove alerts:

- To change the overcommit limit of the storage pool, enter the `alert -set` command.

  ```
  $ alert -set -clustername ClusterA -spname poolA -type overcommit -value 80
  ```

- To view the alert on the storage pool, enter the `alert -list` command.

  ```
  $ alert -list -clustername ClusterA -spname poolA
  ```

  Output similar to the following is displayed:

  ```
  PoolName:          poolA
  PoolID:            FFFFFFFFAC10800E000000004F43B5DA
  ThresholdPercent:  20
  OverCommitPercent: 80
  ```

- To remove the alert on the storage pool, enter the `alert -unset` command.

  ```
  alert -unset -clustername ClusterA -spname poolA -type overcommit
  ```

Related information

alert command

Remove physical volumes from the shared storage pool

On the Virtual I/O Server (VIOS) Version 2.2.3.0, or later, you can remove one or more physical volumes from the shared storage pool by using the command-line interface.

Note: The storage pool must have more than one physical volume. The storage pool must also have free space to accommodate the data of the physical volume that is being removed.

To remove one or more physical volumes from the storage pool:

1. Run the **pv** command. For example,

   ```
   pv -remove -clustername clusterA -sp poolA -pv hdisk2 hdisk3
   ```

   In this example, physical volumes hdisk2 and hdisk3 are removed from the storage pool.
2. Check whether the physical volumes are removed from shared storage pool by using the following command:

   ```
   $ pv -list
   ```

Managing logical units by using the VIOS command line

You can use the command-line interface on the Virtual I/O Server (VIOS) to manage logical units in shared storage pools.

Provisioning client partitions with logical unit storage

You can provision client partitions with logical unit storage by using the command-line interface on the Virtual I/O Server (VIOS).

Creating logical units

You can create logical units and assign the logical units to virtual server adapters by using the Virtual I/O Server (VIOS) command-line interface.

About this task

A logical unit provides the backing storage for the virtual volume of a client partition. By using the following procedure, you can assign a logical unit for each client partition from the shared storage pool of a cluster. Subsequently, you can map the logical unit to the virtual server adapter associated with the virtual Small Computer Serial Interface (SCSI) adapter of the client partition by using the Hardware Management Console (HMC).

When the mapping operations are complete, the logical unit path is similar to the following example:

SAN Storage <=> poolA <=> luA1 <=> viosA1 vtscsi0 <=> viosA1 vhost0 <=> client1 vscsi0 <=> client hdisk0.

Notes:

- A single logical unit can be mapped by multiple virtual server adapters, and thus, accessed by multiple client partitions. However, this mapping typically requires either an additional software layer such as a database management system or the use of the Persistent Reserves standard to manage access to the shared logical unit.
- A logical unit can be mapped from multiple VIOS partitions to a single virtual client.

To create logical units and assign the logical units to virtual server adapters, complete the following steps:

Procedure

1. Obtain the physical location identifiers for the virtual server adapters by running the **`lsmap`** command.

For example, entering the `lsmap -all` command returns results similar to the following:

```
SVSA            Physloc                              Client Partition ID
--------------------------------------------------------------------------
vhost0          U8203.E4A.10D4451-V4-C12             0x00000000

VTD             NO VIRTUAL TARGET DEVICE FOUND
SVSA            Physloc                              Client Partition ID
--------------------------------------------------------------------------
vhost1          U8203.E4A.10D4451-V4-C13             0x00000000
```

Where, `Physloc` identifies the VIOS virtual server adapter related to the HMC property for the `viosA1` VIOS logical partition.

Where:

- `-C12` of the `vhost0` virtual SCSI adapter `physloc` corresponds to the server SCSI adapter ID 12, which maps to virtual SCSI adapter 4 on the `client1` client partition with ID 2
- `-C13` of the `vhost1` virtual SCSI adapter `physloc` corresponds to the server SCSI adapter ID 13, which maps to virtual SCSI adapter 3 on the `client2` client partition with ID 7

The virtual target devices (VTD) also consist of a **Physloc** field. However, the **Physloc** field is empty for VTDs because the HMC property is not applicable to a VTD.

2. Create the logical unit by running the **`mkbdsp`** command.

   For example:

   - The `luA1` logical unit is created in the `poolA` storage pool of the `clusterA` cluster, with thin-provisioning and an initial provisional size of 100 MB.

     ```
     mkbdsp -clustername clusterA -sp poolA 100M -bd luA1
     ```

   - The `luA3` logical unit is created in the `poolA` storage pool of the `clusterA` cluster, with thick-provisioning and an initial provisional size of 100 MB.

     ```
     mkbdsp -clustername clusterA -sp poolA 100M -bd luA3 -thick
     ```

3. Map the logical unit to the virtual server adapter associated with the client partition by running the **`mkbdsp`** command.

   For example:

   - The `luA1` logical unit is thin-provisioned and mapped to the `vscsi0` virtual server adapter associated with the `client1` client partition, which the HMC properties and the **`lsmap`** command indicate as `vhost0`.

     ```
     mkbdsp -clustername clusterA -sp poolA -bd luA1 -vadapter vhost0
     ```

   - The `luA3` logical unit is thick-provisioned and mapped to the `vscsi0` virtual server adapter associated with the `client1` client partition, which the HMC properties and the **`lsmap`** command indicate as `vhost0`.

     ```
     mkbdsp -clustername clusterA -sp poolA -bd luA3 -vadapter vhost0 -thick
     ```

4. Create the logical unit in the shared storage pool, and map it to the virtual server adapter associated with the client partition.

   For example:

   - The `luA2` logical unit is created in the `poolA` storage pool of the `clusterA` cluster, with thin-provisioning and an initial provisional size of 200 MB. The `luA2` logical unit is then mapped to the `vscsi0` virtual server adapter associated with the `client2` client partition, which the HMC properties and the **`lsmap`** command indicate as `vhost1`.

     ```
     mkbdsp -clustername clusterA -sp poolA 200M -bd luA2 -vadapter vhost1 -tn vtscsi1
     ```

   - The `luA4` logical unit is created in the `poolA` storage pool of the `clusterA` cluster, with thick-provisioning and an initial provisional size of 200 MB. The `luA4` logical unit is then mapped to the vscsi0 virtual server adapter associated with the client2 client partition, which the HMC properties and the **lsmap** command indicate as vhost1.

```
mkbdsp -clustername clusterA -sp poolA 200M -bd luA4 -vadapter vhost1 -tn vtscsi1 -thick
```

Note: The -tn vtscsiX option is not mandatory. If this option is omitted, a default value is used. By specifying the virtual target name, you can run the **lsdevinfo** command and search for information by using the target name. In addition, you can map multiple logical units to the same virtual host adapter. The virtual target name is used to distinguish the mappings.

5. Display the logical unit information.

   For example, entering the lssp -clustername clusterA -sp poolA -bd command returns results similar to the following. Here, the logical unit is the backing device, or bd.

```
LU Name   Size (MB)    ProvisionType   %Used Unused(mb) LU UDID
-------------------------------------------------------------------------------------------
luA1      100          THIN            10%   90         258f9b298bc302d9c7ee368ff50d04e3
luA2      200          THIN            15%   170        7957267e7f0ae3fc8b9768edf061d2f8
luA3      100          THICK           5%    95         459f9b298bc302fc9c7ee368f50d04e3
luA4      200          THICK           0%    200        6657267e7d0ae3fc7b9768edf061d2d2
```

   Entering the lsmap -all command returns results similar to the following:

```
SVSA            Physloc                                      Client Partition ID
----------------------------------------------------------------------------------------
vhost0          U8203.E4A.10D4451-V4-C12                     0x00000002

VTD                   vtscsi0
Status                Available
LUN                   0x8100000000000000
Backing device        lua1.b1277fffdd5f38acb365413b55e51638
Physloc
Mirrored              N/A

VTD                   vtscsi1
Status                Available
LUN                   0x8200000000000000
Backing device        lua2.8f5a2c27dce01bf443383a01c7f723d0
Physloc
Mirrored              N/A
```

Results

In the examples in this topic, the vscsi0 virtual client SCSI adapter on client partitions Client1 and Client2 was mapped to the logical units luA1 and luA2.

Related information

lsmap command

lssp command

mkbdsp command

Enabling the logical unit backed storage
You can enable the logical unit backed storage by using the Virtual I/O Server (VIOS) command-line interface.

About this task

To display the virtual physical volumes that the logical units represent in the client environments and enable the logical unit backed storage, complete the following steps:

Procedure

1. Log in to the client as root user.
2. Enter the following commands in the Korn shell:

```
cfgmgr
lspv
lsdev -c adapter -F 'name physloc'
lsdev -t vdisk  -F 'name physloc'
```

The **cfgmgr** command reassembles device configuration information and picks up the new mapping for the virtual Small Computer Serial Interface (SCSI) adapter. The **lspv** and **lsdev** commands on the client, and the **lsdev** command on the VIOS can be used to verify the association of the *hdiskX* physical volume and the *vscsiX* virtual SCSI adapter to the *vhostY* virtual server adapter on the VIOS partition (where, X and Y are appropriate instance numbers). After the vscsi*X* to hdisk*X* mapping is verified, the normal volume group, file system management, and I/O can proceed on the client partitions, as if the *hdiskX* physical volume is another direct connection SCSI device. Other than establishing the client virtual physical volume association with a VIOS path, no further action is required on the client. Hence, you can exit the client shell.

These steps are unique to the AIX client. The Linux operating system also supports adding new storage devices dynamically. Run the following commands:

```
ls -vscsi
lsscsi
echo "- - -" > /sys/class/scsi_host/hostX/scan
lsscsi
cat /sys/class/scsi_host/hostX/partition_name
```

The **ls -vscsi** command displays all virtual SCSI host adapters. The **partition_name** attribute displays the connected VIOS partition. Replace *hostX* with the host number to which storage has been added. The **lsscsi** command displays all attached SCSI disks.

Note: When new data is written to the *hdiskX* physical volume, the VIOS logical partition monitors for overruns of threshold limits. A shell connection to each of the logical partitions must be maintained to observe threshold alerts in the VIOS error log. Alerts can also be captured by using management tools. The threshold limit can be changed to avoid or delay the alerts.

3. Enter the following commands:

```
ls -vscsi
lsscsi
echo "- - -" > /sys/class/scsi_host/hostX/scan
lsscsi
cat /sys/class/scsi_host/hostX/partition_name
```

The **ls -vscsi** command displays all virtual SCSI host adapters. The **partition_name** attribute displays the connected VIOS partition. Replace *hostX* with the host number to which storage has been added. The **lsscsi** command displays all attached SCSI disks.

Note: When new data is written to the *hdiskX* physical volume, the VIOS logical partition monitors for overruns of threshold limits. A shell connection to each of the logical partitions must be maintained to observe threshold alerts in the VIOS error log. Alerts can also be captured by using management tools. The threshold limit can be changed to avoid or delay the alerts.

Related information

cfgmgr command

lsdev command

lspv command

Increasing the size of an existing logical unit

You can use the command-line interface on the Virtual I/O Server (VIOS) to increase the size (resize) of an existing logical unit (LU).

Before you begin

You can use the resize feature to increase the size of the existing LUs.

An LU can be thick or thin provisioned. You can change the size of both thick and thin provisioned LUs. You can also change the size of an LU while it is mapped to one or more clients, and I/O can be occurring to the LU at the time.

To increase the size of an LU, the LU must be uniquely identified by the name or UDID.

About this task

You can increase the size of an LU by completing this example procedure:

Procedure

Enter the following command: `lu -resize -lu` *`luName`* `-size` *`newSize`*

Results

The size of the named LU is increased to the *newSize* you specified.

The resize LU operation and snapshots

Consider the following sequence of events:

1. You take a snapshot of an LU.
2. You perform the resize operation to increase the capacity of the LU.
3. Then you perform a rollback of the LU to the previous snapshot.

The rollback task changes the size of the LU back to the original state at the time of the snapshot. This is effectively a decrease in the capacity of the LU, which is not supported. To prevent such a scenario, the resize LU operation determines whether the LU has any snapshots for rollback purposes. If it finds such snapshots, resize LU fails with an appropriate exception message.

Note: Snapshots that are created for cloning are not relevant, and the resize operation succeeds if only cloning snapshots are present.

Limitations for resizing LUs

- A single resize LU operation does not support multiple LUs. This means that to perform a resize operation on multiple LUs, you need to make multiple resize LU requests.
- The following operations are mutually exclusive with a resize operation, which means that when one of the operations is occurring on a particular LU, you cannot resize that LU:
  - Remove LU
  - Map LU
  - Unmap LU
  - Initialize LU
  - Sanpshot
  - Rollback
  - Move LU
  - Live Partition Mobility (LPM) of a client mapped to the LU
  - Another resize LU
- If the new capacity of the LU that you provide is smaller than the current capacity of the LU, the operation fails.

Moving a logical unit from one storage tier to another

A logical unit (LU) can be moved from one storage tier to another storage tier. One LU, tree, or subtree can be moved at a time.

Before you begin

When a logical unit (LU) shares storage blocks with other LUs, it is part of a logical subtree. LU subtrees can exist when you use a management tool, such as IBM Power Virtualization Center (PowerVC) to deploy clients. LU subtrees cannot be created from the VIOS command line interface. You can move any kind of LU. LUs can have further snapshots and clones. Clones are based on snapshots and hence, clones inherit blocks from the snapshot. An LU within a subtree is categorized as one of the following node types:

root
: This is the first level of the subtree. This LU is a parent LU to all of the other LUs in the tree.

intermediate
: This is a middle level of the subtree, and has at least one parent LU and at least one child LU in the subtree.

leaf
: This is the final level of the subtree. LUs at this level must have a parent LU, but no child LUs. If it has no parent LU and no child LUs, then it is a root LU.

About this task

When you move an LU that is part of a subtree, all of the children of that LU also move. The subtree can be broken by using the `-nonrecursive` flag in the command. LU move is tracked by using the LU_MOVE_STATUS in `lu –list` output.

You can move an LU from one storage tier to another storage tier, by completing this example procedure:

Procedure

Enter the following command: `lu -move -lu` *`luName`* `-dsttier` *`newTier`*.

For this example step, any children in the subtree are also moved to the new storage tier. If you want to break the relationship and not move the child LUs, use the **`-nonrecursive`** parameter in the command. When you use the **`-nonrecursive`** parameter, overall disk usage increases because blocks shared with other LUs are no longer shared with the LU that has been moved.

Results

An LU belongs to only one storage tier at a time, which is called the *primary* storage tier. During the move, the destination storage tier is the *primary* tier. The *primary* storage tier is set before moving data blocks. The data blocks are moved in the background. During the move, the LU resides in multiple storage tiers, with some blocks in the destination storage tier and some blocks in the source storage tier.

Logical unit move failures

A common cause of an LU failure during a move is a lack of space on the destination storage tier. If a move fails, the LU remains in a failed condition and the LU has blocks in both the source storage tier and in the destination storage tier. To recover a failed LU move, you must clear the existing LUs or add new PVs to the destination tier and restart the move. The LU continues to operate normally in this state, so there is no interruption in access to the LU. In such a scenario, manual intervention is required to recover from the failure.

Listing the storage tiers of a logical unit

Working with a logical unit (LU) requires that you identify which storage tiers contain blocks for this LU.

About this task

To list the tiers that contain blocks of a certain LU, enter the following command:

Procedure

`lu -list -verbose`.

The following information is provided that helps you identify the tier relationships:

TIER_NAME
The name of the tier that the information applies to.

TIER_RELATION
The status of the listed tier for the LU. Values are PRIMARY (the destination or only storage tier) or VACATING (a source tier in a failed move). If the value is vacating, another storage tier is related to this LU.

ADDITIONAL_TIERS
Other storage tiers that contain blocks of this LU.

LU_MOVE_STATUS
The last known status of a move for this LU. Values can be: N/A, in progress, failed, recursive success, recursive in progress, or recursive failed.

Unmapping a logical unit

You can unmap a logical unit by using the Virtual I/O Server (VIOS) command-line interface.

About this task

To unmap logical units from the virtual server adapter, complete the following steps:

Procedure

1. To display the mapping of the virtual server adapter, enter the `lsmap -all` command.

```
SVSA            Physloc                                      Client Partition ID
--------------- -------------------------------------------- ------------------
vhost0          U8203.E4A.10D4451-V4-C12                     0x00000002

VTD                   vtscsi0
Status                Available
LUN                   0x8100000000000000
Backing device        testLU.b1277fffdd5f38acb365413b55e51638
Physloc
Mirrored              N/A

VTD                   vtscsi1
Status                Available
LUN                   0x8200000000000000
Backing device        test_LU.8f5a2c27dce01bf443383a01c7f723d0
Physloc
Mirrored              N/A
```

2. To unmap a logical unit, run the **rmbdsp** command with the **-vtd** option. If you do not use the **-vtd** option, the entire logical unit is removed.

   In the following example, the mapping for the *luA2* logical unit is removed.

```
rmbdsp -vtd vtscsi1
```

Related information

lsmap command

rmbdsp command

Removing logical units

You can remove logical units from the shared storage pool by using the Virtual I/O Server (VIOS) command-line interface.

Before you remove the logical units from shared storage pools, you must delete the mapping of physical volumes by reconfiguring the clients that reference the logical unit path.

To remove a logical unit from the shared storage pool, use the following commands, as applicable:

- To display logical unit information, run the **lssp** command. For example, entering the `lssp -clustername clusterA -sp poolA -bd` command returns results similar to the following:

  ```
  LU Name   Size (MB)    ProvisionType   %Used Unused(mb) LU UDID
  ---------------------------------------------------------------------------------------
  luA1      100          THIN            10%   90         258f9b298bc302d9c7ee368ff50d04e3
  luA2      200          THIN            15%   170        7957267e7f0ae3fc8b9768edf061d2f8
  luA3      100          THICK           5%    95         459f9b298bc302fc9c7ee368f50d04e3
  luA4      200          THICK           0%    200        6657267e7d0ae3fc7b9768edf061d2d2
  ```

- To remove a logical unit, run the **rmbdsp** command. For example:

  ```
  rmbdsp -clustername clusterA -sp poolA -bd luA2
  ```

  Notes:

  – Returning a logical unit to the shared storage pool might cause a storage threshold transition alert.

  – If the logical unit is still mapped to a different VIOS logical partition, the **rmbdsp** command fails.

  – If the logical unit is only mapped to virtual server adapters on the same VIOS logical partition on which you run the command, the mappings and the logical unit are deleted. To see the VIOS logical partition that actually has the logical unit mapped, run the **lsmap -clustername** command.

- To remove one of the multiple logical units with the same name, specify the unique identifier of the logical unit. For example, when there is a second logical unit `luA1`, entering the following command removes that logical unit.

  ```
  rmbdsp -clustername clusterA -sp poolA -luudid 258f9b298bc302d9c7ee368ff50d04e3
  ```

- To remove all the logical units from the shared storage pool, run the **rmbdsp** command with the **-all** option.

  ```
  rmbdsp -clustername clusterA -sp poolA -all
  ```

  The shared storage pool is not removed when all the logical units are removed. All physical volumes previously added to the pool remain in the pool and cannot be removed when the pool exists. Delete the cluster to delete the default pool and recover the physical volumes.

  To remove all the logical units, there must be no virtual target device assigned to any logical unit. Run the **rmbdsp** command with the **-vtd** option on each virtual target device assigned to the logical units to ensure that no virtual target device is assigned to any logical unit.

Related tasks

Deleting a cluster
You can delete a cluster by using the Virtual I/O Server (VIOS) command-line interface.

Related information

lssp command

rmbdsp command

Migrating a cluster configuration

You can migrate the cluster that you created and configured on the VIOS logical partition that has Version 2.2.0.11, Fix Pack 24, Service Pack 1 to the VIOS logical partition that has Version 2.2.1.0. By performing this task, you can restore the previous shared storage pool mappings with a new shared storage pool and database versions.

About this task

To migrate a cluster that you created and configured on the VIOS logical partition that has Version 2.2.0.11, Fix Pack 24, Service Pack 1 to the VIOS logical partition that has Version 2.2.1.0, or later:

Procedure

1. Create a backup of the cluster that you want to migrate on the VIOS logical partition that has Version 2.2.0.11, Fix Pack 24, Service Pack 1. For example:

   ```
   viosbr -backup -file oldCluster -clustername clusterA
   ```

   Save the backup file that is generated on a different system. For example: `oldCluster.clusterA.tar.gz`.
2. Reinstall the VIOS logical partition that has Version 2.2.1.0, or later.

   Note: Do not change the physical volumes that are used for the storage pool.
3. Migrate the backup file that was created in step 1 to the VIOS logical partition that has Version 2.2.1.0, or later. For example:

   ```
   viosbr -migrate -file oldCluster.clusterA.tar.gz
   ```

   This step migrates the backup file to the VIOS logical partition with VIOS Version 2.2.1.0, or later. For example: `oldCluster_MIGRATED.clusterA.tar.gz`.
4. Clean the physical volume that is to used as the cluster repository disk. For example:

   ```
   cleandisk -r hdisk9
   ```

   Note: Do not change the physical volumes that are used for the storage pool.
5. Restore the network devices by using the migrated backup file. For example:

   -
     ```
     viosbr -restore -file oldCluster_MIGRATED.clusterA.tar.gz -clustername clusterA -repopvs
     hdisk9
     -type net
     ```

   -
     ```
     viosbr -restore -file oldCluster_MIGRATED.clusterA.tar.gz -clustername clusterA -subfile
     clusterAMTM9117-MMA0206AB272P9.xml -type net
     ```

   Note: With VIOS Version 2.2.2.0, and later, you do not need to restore the network devices before restoring a cluster when you are migrating a cluster configuration. Hence, if you are using VIOS Version 2.2.2.0, and later, skip this step.
6. Restore the cluster by using the migrated backup file. For example:

   -
     ```
     viosbr -restore -file oldCluster_MIGRATED.clusterA.tar.gz -clustername clusterA -repopvs
     hdisk9
     ```

   -
     ```
     viosbr -restore -file oldCluster_MIGRATED.clusterA.tar.gz -clustername clusterA -subfile
     clusterAMTM9117-MMA0206AB272P9.xml
     ```

   After a successful restore operation, the cluster and all shared storage pool mappings are configured as in the VIOS logical partition that has Version 2.2.0.11, Fix Pack 24, Service Pack 1.
7. Verify that the cluster restored successfully by listing the status of the nodes in the cluster. For example:

   ```
   cluster -status -clustername clusterA
   ```
8. List the storage mappings on the VIOS. For example:

   ```
   lsmap -all
   ```

   Note: To migrate a cluster from VIOS Version 2.2.1.3 to VIOS Version 2.2.2.0, ensure that the rolling update procedure is completed.

Related concepts

Rolling updates in a cluster

The Virtual I/O Server (VIOS) Version 2.2.2.0 supports rolling updates for clusters.

Rolling updates in a cluster

The Virtual I/O Server (VIOS) Version 2.2.2.0 supports rolling updates for clusters.

You can use the rolling updates enhancement to apply software updates to the VIOS logical partitions in the cluster individually without causing an outage in the entire cluster. The updated logical partitions cannot use the new capabilities until all logical partitions in the cluster are updated and the cluster is upgraded.

To update the VIOS logical partitions to use the new capabilities, ensure that the following conditions are met:

- All VIOS logical partitions must have the new level of software installed. You can verify that the logical partitions have the new level of software installed by typing the `cluster -status -verbose` command from the VIOS command line. In the `Node Upgrade Status` field, if the status of the VIOS logical partition is displayed as `UP_LEVEL`, the software level in the logical partition is higher than the software level in the cluster. If the status is displayed as `ON_LEVEL`, the software level in the logical partition and the cluster is the same.
- All VIOS logical partitions must be running. If any VIOS logical partition in the cluster is not running, the cluster cannot be upgraded to use the new capabilities.

The VIOS logical partition that is acting as the database primary node (DBN) periodically checks whether an upgrade is required. This check is done in 10-minute intervals. Only the DBN is allowed to initiate and coordinate an upgrade.

Restrictions: When an upgrade is being performed, the following cluster configuration operations are restricted:

- Adding a VIOS logical partition to the cluster
- Adding a physical volume to the shared storage pool
- Replacing a physical volume in the shared storage pool
- Removing physical volumes from the shared storage pool

Getting started with shared storage pools by using the VIOS configuration menu

Learn about using the Virtual I/O Server (VIOS) configuration menu to manage shared storage pools.

On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, you can create a clustering configuration. The VIOS partition in a cluster is connected to the shared storage pool. VIOS partitions that are connected to the same shared storage pool must be part of the same cluster. Each cluster has a default storage pool. You can use the VIOS command-line interface to manage shared storage pools.

Notes:

- On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, a cluster consists of only one VIOS partition.
- VIOS Version 2.2.1.0 supports only one cluster in a VIOS partition.
- On VIOS Version 2.2.1.3, or later, a cluster consists of up to four networked VIOS partitions.
- On VIOS Version 2.2.2.0, or later, a cluster consists of up to 16 networked VIOS partitions.

To access VIOS configuration menu, run the **`cfgassist`** command from the command-line interface. On the VIOS configuration menu, move the cursor to the **Shared Storage Pools** menu and press Enter. Use the submenus to manage clusters, VIOS logical partitions, storage pools, and logical units in shared storage pools.

To select information, such as existing cluster names, associated storage pool names, snapshot names, and logical unit names on the **Shared Storage Pools** menu, you can use the following wizards on the VIOS configuration menu:

- Cluster and Storage Pool Selection wizard: On the **Shared Storage Pools** menu, you can use the Cluster and Storage Pool Selection wizard to select the name of an existing cluster and associated storage pool. The Cluster and Storage Pool Selection wizard displays the set of cluster names. After you select a cluster, the wizard displays the names of the associated storage pools.
- Logical Unit Selection wizard: On the **Manage Logical Units in Storage Pool** submenu, you can use the Logical Unit Selection wizard to select the names of logical units. You can identify multiple logical unit names, redisplay the Logical Unit Selection wizard, and change the logical unit selection.
- Snapshot Selection wizard: On the **Manage Logical Units in Storage Pool** submenu, you can use the Snapshot Selection wizard to select snapshots and logical units. You can select cluster names and the storage pool name.

Related information

cfgassist command

Managing a cluster by using the VIOS configuration menu

You can use the Virtual I/O Server (VIOS) configuration menu to manage a cluster and the Virtual I/O Server logical partitions.

Creating a cluster

You can create a cluster in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To create a cluster in shared storage pools:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Cluster and VIOS Nodes** submenu and press Enter.
2. From the **Manage Cluster and VIOS Nodes** submenu, move the cursor to the **Create Cluster** option and press Enter.

   The Create Cluster window opens.
3. Enter the cluster name being created in the **Cluster name** field.
4. Enter the storage pool name in the **Storage Pool name** field.
5. Press F4 or Esc + 4 in the **Physical Volumes for Repository** field and select the repository physical volumes.
6. Press F4 or Esc + 4 in the **Physical Volumes for Storage Pool** field and select the physical volume names for the storage pool.
7. To clean the physical volumes, type `yes` in the **Clean physical volumes before use** field. Otherwise, type `no`.
8. Press Enter to create a cluster.
9. In the confirmation window that opens, select **Yes** to proceed with creating the cluster.

Listing all clusters

You can list all clusters in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list all clusters in shared storage pools:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Cluster and VIOS Nodes** submenu and press Enter.
2. From the **Manage Cluster and VIOS Nodes** submenu, move the cursor to the **List All Clusters** option and press Enter.

   The list of all clusters that are associated with the VIOS logical partition is displayed.

Deleting a cluster

You can delete a cluster from shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

Notes:

- You cannot restore a cluster if you delete the cluster. You cannot restore a VIOS logical partition in a cluster if the VIOS logical partition is removed from the cluster.
- Deleting a cluster fails if the VIOS logical partition has any mappings to logical units in the shared storage pool or if there are any logical units within the shared storage pool. Before you perform the delete operation, remove all logical partition mappings and logical units.

To delete a cluster from shared storage pools:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Cluster and VIOS Nodes** submenu and press Enter.
2. From the **Manage Cluster and VIOS Nodes** submenu, move the cursor to the **Delete Cluster** option and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name to be deleted.

   The Delete Cluster window displays the cluster name that you selected.
4. Press Enter to delete the cluster.
5. In the confirmation window that opens, select **Yes** to proceed with deleting the cluster.

Related concepts

Unmapping logical units
Learn about unmapping logical units by using the Virtual I/O Server (VIOS) configuration menu.

Related tasks

Deleting a logical unit
You can delete a logical unit from shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

Adding VIOS nodes to a cluster

You can add Virtual I/O Server (VIOS) nodes to a cluster by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To add VIOS nodes to a cluster:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Cluster and VIOS Nodes** submenu and press Enter.

2. From the **Manage Cluster and VIOS Nodes** submenu, move the cursor to the **Add VIOS Nodes to Cluster** option and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name.

   The Add VIOS Nodes to Cluster window displays the cluster name that you selected.
4. Enter the VIOS node names in the **Network names of Nodes to add** field. Enter multiple node names separated by a space.
5. Press Enter to add the VIOS nodes.
6. In the confirmation window that opens, select **Yes** to proceed with adding the VIOS nodes.

Deleting VIOS nodes from a cluster

You can delete Virtual I/O Server (VIOS) nodes from a cluster by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To delete VIOS nodes from a cluster:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Cluster and VIOS Nodes** submenu and press Enter.
2. From the **Manage Cluster and VIOS Nodes** submenu, move the cursor to the **Delete Nodes from Cluster** option and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name.

   The nodes of the cluster are displayed.
4. Select one or more nodes and press Enter.

   The Delete VIOS Nodes From Cluster window opens.
5. Press F4 or Esc + 4 in the **Network names of Nodes to delete** field to change the node selection.
6. Press Enter to delete the VIOS nodes.
7. In the confirmation window that opens, select **Yes** to proceed with deleting the VIOS nodes.

   Note: If the VIOS logical partition is mapped to a logical unit in the storage pool of the cluster, deleting VIOS nodes from a cluster fails. To remove the logical partition, unmap the logical unit.

Related concepts

Unmapping logical units
Learn about unmapping logical units by using the Virtual I/O Server (VIOS) configuration menu.

Listing VIOS nodes in a cluster

You can list all Virtual I/O Server (VIOS) nodes in a cluster by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list all Virtual I/O Server nodes in a cluster:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Cluster and VIOS Nodes** submenu and press Enter.
2. From the **Manage Cluster and VIOS Nodes** submenu, move the cursor to the **List Nodes in Cluster** option and press Enter.

3. Select the cluster name in the window that opens.

   The list of all VIOS nodes associated with the cluster is displayed.

Managing storage pools by using the VIOS configuration menu

You can use the Virtual I/O Server (VIOS) configuration menu to manage shared storage pools.

Listing storage pools in a cluster

You can list storage pools in a cluster by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list storage pools in a cluster:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Storage Pools in Cluster** submenu and press Enter.
2. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **List Storage Pools in Cluster** option, and press Enter.
3. Select the cluster name in the window that opens.

   The list of all storage pools that are associated with the cluster is displayed.

Listing physical volumes in the storage pool

You can list the physical volumes in the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list the physical volumes in the storage pool:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Storage Pools in Cluster** submenu and press Enter.
2. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **List Physical Volumes in Storage Pool** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The list of all physical volumes associated with the storage pool is displayed.

Adding storage space to the storage pool

When more storage space is required in a storage pool, you can use the Virtual I/O Server (VIOS) configuration menu to add one or more physical volumes to the storage pool.

Adding physical volumes to the storage pool

You can add physical volumes to the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To add physical volumes to the storage pool:

Procedure

1. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **Change/Show Physical Volumes in Storage Pool** submenu, and press Enter.
2. From the **Change/Show Physical Volumes in Storage Pool** submenu, move the cursor to the **Add Physical Volumes to Storage Pool** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Add Physical Volumes to Storage Pool** window displays the cluster name and storage pool name that you selected.
5. Press F4 or Esc + 4 in the **Physical Volumes to add** field and select the physical volume. You can select multiple physical volumes.
6. To clean the physical volumes, type `yes` in the **Clean physical volumes before use** field. Otherwise, type `no`.
7. Press Enter to add the physical volumes to the storage pool.
8. In the confirmation window that opens, select **Yes** to proceed with adding the physical volumes to the storage pool.

Replacing physical volumes in the storage pool
You can replace physical volumes in the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

When more storage space is needed in a storage pool, you can also remove and replace existing physical volumes in a storage pool. If you are replacing the existing physical volume with a physical volume that has a larger capacity, the capacity of the shared storage pool increases.

Restrictions:

- You can replace physical volumes only in one cluster at a time.
- Do not use this task to increase only the capacity of the shared storage pool.

To remove and replace physical volumes in the storage pool:

Procedure

1. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **Change/Show Physical Volumes in Storage Pool** submenu and press Enter.
2. From the **Change/Show Physical Volumes in Storage Pool** submenu, move the cursor to the **Replace Physical Volumes in Storage Pool** option and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Replace Physical Volumes in Storage Pool** window displays the cluster name and storage pool name that you selected.
5. Press F4 or Esc + 4 in the **Physical Volumes to replace** field and select the physical volume. You can select multiple physical volumes.
6. Press F4 or Esc + 4 in the **Physical Volumes to add** field and select the physical volume. You can select multiple physical volumes.
7. Press Enter to replace the physical volumes in the storage pool.
8. In the confirmation window that opens, select **Yes** to proceed with replacing the physical volumes in the storage pool.

Results

Note: If the size of physical volume that is being replaced is large, the replace operation might take a longer time to complete.

Listing physical volumes in the storage pool
You can list the physical volumes in the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list the physical volumes in the storage pool:

Procedure

1. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **Change/Show Physical Volumes in Storage Pool** submenu and press Enter.
2. From the **Change/Show Physical Volumes in Storage Pool** submenu, move the cursor to the **Physical Volumes** option and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The list of all physical volumes that are associated with the storage pool is displayed.

Setting and modifying the storage pool threshold alert

You can use the Virtual I/O Server (VIOS) configuration menu to perform tasks that are related to setting or modifying the storage pool threshold alert on the VIOS configuration menu.

Listing the storage pool threshold alert value
You can list the threshold alert value of the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list the threshold alert value of the storage pool:

Procedure

1. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **Set/Modify Storage Pool Threshold Alert** submenu and press Enter.
2. From the **Set/Modify Storage Pool Threshold Alert** submenu, move the cursor to the **List threshold alert levels in Storage Pool** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The threshold alert value of the storage pool is displayed.

Changing the storage pool threshold alert value
You can change the threshold alert value of the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To change the threshold alert value of the storage pool:

Procedure

1. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **Set/Modify Storage Pool Threshold Alert** submenu and press Enter.
2. From the **Set/Modify Storage Pool Threshold Alert** submenu, move the cursor to the **Change threshold alert level in Storage Pool** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Change Threshold Alert Level in Storage Pool** window displays the cluster name, storage pool name, and current threshold alert value of the storage pool.
5. Enter the new threshold alert value in the **New threshold alert level** field.
6. Press Enter to update the new threshold alert value.

Removing the storage pool threshold alert value
You can remove the threshold alert value of the storage pool by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To remove the threshold alert value of the storage pool:

Procedure

1. From the **Manage Storage Pools in Cluster** submenu, move the cursor to the **Set/Modify Storage Pool Threshold Alert** submenu and press Enter.
2. From the **Set/Modify Storage Pool Threshold Alert** submenu, move the cursor to the **Remove threshold alert level in Storage Pool** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Remove Threshold Alert Level in Storage Pool** window displays the cluster name and storage pool name that you selected.
5. Press Enter to remove the threshold alert level of the storage pool.

Managing logical units by using the VIOS configuration menu

You can use the Virtual I/O Server (VIOS) configuration menu to manage logical units in shared storage pools.

Creating and mapping logical units

You can create and map logical units in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To create and map logical units in shared storage pools:

Procedure

1. From the **Shared Storage Pools** menu, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Create and Map Logical Unit** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.

3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Create and Map Logical Unit** window displays the cluster name and storage pool name that you selected.
5. Enter the logical unit name being created in the **Logical Unit name** field.
6. Enter the logical unit size in megabytes in the **Logical Unit size** field.
7. Press F4 or Esc + 4 in the **Virtual server adapter to map** field and select the virtual server adapter name that you want to map.
8. Enter the virtual target device name in the **Virtual target device name** field.
9. Press Enter to create and map the logical unit.

Creating logical units

You can create logical units in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To create logical units in shared storage pools:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Create Logical Unit** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Create Logical Uni**t window displays the cluster name and storage pool name that you selected.
5. Enter the logical unit name being created in the **Logical Unit name** field.
6. Enter the logical unit size in megabytes in the **Logical Unit size** field.
7. Press Enter to create the logical unit.

Mapping logical units

You can map an existing logical unit to a virtual server adapter in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To map an existing logical unit to a virtual server adapter in shared storage pools:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Map Logical Unit** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The Logical Unit Selection wizard starts.
5. Select the logical unit name and press Enter.

The **Map Logical Unit** window displays the cluster name, storage pool name, and the logical unit name that you selected.

6. Press F4 or Esc + 4 in the **Virtual server adapter to map** field and select the virtual server adapter name that you want to map.
7. Enter the virtual target device name in the **Virtual target device name** field.
8. Press Enter to map the logical unit.

Unmapping logical units

Learn about unmapping logical units by using the Virtual I/O Server (VIOS) configuration menu.

Unmapping logical units by logical unit name

You can unmap logical units by selecting the logical unit names by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To unmap logical units by selecting the logical unit names:

Procedure

1. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Unmap Logical Unit** submenu and press Enter.
2. From the **Unmap Logical Unit** submenu, move the cursor to the **Unmap Logical Unit by LU Name** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Logical Unit Selection By LU Name** window opens.
5. Move the cursor to the logical unit names that you want to unmap and press F7 (function key 7). You can select multiple logical unit names. To unmap all logical units, select **ALL**.
6. Press Enter after you select the logical units to unmap.

   The **Unmap Logical Unit By LU Name** window displays the cluster name, storage pool name, and the logical unit names that you selected.
7. Type `yes` in the **Confirm unmap** field to confirm that you want to unmap the logical units.
8. Press Enter to unmap the logical units.

Unmapping logical units by virtual server adapter name

You can unmap logical units by virtual server adapter name by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To unmap logical units by selecting the virtual server adapter names:

Procedure

1. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Unmap Logical Unit** submenu and press Enter.
2. From the **Unmap Logical Unit** submenu, move the cursor to the **Unmap Logical Unit by Virtual Server Adapter Name** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

The **Logical Unit Selection By Virtual Server Adapter Name** window opens.

5. Move the cursor to the virtual server adapter names corresponding to the logical unit that you want to unmap and press F7 (function key 7). You can select multiple virtual server adapter names. To select all virtual server adapter names, select **ALL**.
6. Press Enter after you select the virtual server adapter names.

   The **Unmap Logical Unit By VAdapter** window displays the cluster name, storage pool name, and the logical unit names corresponding to the virtual server adapter names that you selected.
7. Type yes in the **Confirm unmap** field to confirm that you want to unmap the logical units.
8. Press Enter to unmap the logical units.

Unmapping logical units by virtual target device name
You can unmap logical units by virtual target device name by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To unmap logical units by selecting the virtual target device names:

Procedure

1. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Unmap Logical Unit** submenu and press Enter.
2. From the **Unmap Logical Unit** submenu, move the cursor to **Unmap Logical Unit by Virtual Target Device Name**, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The **Logical Unit Selection By Virtual Target Device Name** window opens.
5. Move the cursor to the virtual target device names corresponding to the logical unit that you want to unmap and press F7 (function key 7). You can select multiple virtual target device names. To select all virtual target device names, select **ALL**.
6. Press Enter after you select the virtual target device names.

   The **Unmap Logical Unit By VTD** window displays the cluster name, storage pool name, and the logical unit names corresponding to the virtual target device names that you selected.
7. Type yes in the **Confirm unmap** field to confirm that you want to unmap the logical units.
8. Press Enter to unmap the logical units.

Deleting a logical unit

You can delete a logical unit from shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To delete a logical unit from shared storage pools:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Delete Logical Unit** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.

4. Select the storage pool name and press Enter.

   The Logical Unit Selection wizard starts.
5. Select the logical unit name and press Enter.

   The **Delete Logical Unit** window displays the cluster name, storage pool name, and the logical unit name that you selected.
6. Press Enter to delete the logical unit.
7. In the confirmation window that opens, select **Yes** to proceed with deleting the logical unit.

Listing logical units

You can list logical units in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list logical units in shared storage pools:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **List Logical Units** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The list of all logical units that are associated with the shared storage pool is displayed.

Listing logical unit maps

You can list the logical unit mappings in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list the logical unit mappings in shared storage pools:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **List Logical Unit Maps** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The list of all logical unit mappings that are associated with the shared storage pool is displayed.

Creating a logical unit snapshot

You can create snapshots of logical units in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu. Snapshots are images of a single logical unit or multiple logical units.

Before you begin

Note: Before you create a snapshot, perform synchronization of the virtual disk on the client partition.

About this task

To create snapshots of logical units in shared storage pools:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Create Logical Unit Snapshot** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.

   The Logical Unit Selection wizard starts.
5. Select the logical unit names and press Enter.

   The **Create Logical Unit Snapshot** window displays the cluster name, storage pool name, and the logical unit names that you selected.
6. Enter the snapshot name in the **Snapshot name** field.
7. Press Enter to create the snapshot of the logical units.

Listing logical unit snapshots

Learn about listing snapshots of logical units by using the Virtual I/O Server (VIOS) configuration menu. Snapshots are images of a single logical unit or multiple logical units.

Listing snapshots for a logical unit

You can list snapshots for a logical unit in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list snapshots for a logical unit in shared storage pools:

Procedure

1. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **List Logical Unit Snapshot** submenu and press Enter.
2. From the **List Logical Unit Snapshot** submenu, move the cursor to the **List Snapshots for a Logical Unit** option and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.
5. Select the logical unit name in the window that opens and press Enter.

   The **List Snapshots for a Logical Unit** window displays the cluster name, storage pool name, and the logical unit names.
6. Press Enter to display the set of snapshots that are associated with the selected logical unit.

Listing logical units in a snapshot

You can list the logical units in a snapshot in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list the logical units in a snapshot:

Procedure

1. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **List Logical Unit Snapshot** submenu, and press Enter.
2. From the **List Logical Unit Snapshot** submenu, move the cursor to the **List Logical Units in a Snapshot** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.
5. Select the snapshot name in the window that opens.

   The **List Logical Units in a Snapshot** window displays the cluster name, storage pool name, and the snapshot name.
6. Press Enter to display the set of logical units that are associated with the selected snapshot.

Listing all logical unit snapshots
You can list all logical unit snapshots in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu.

About this task

To list all logical unit snapshots in shared storage pools:

Procedure

1. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **List Logical Unit Snapshot** submenu, and press Enter.
2. From the **List Logical Unit Snapshot** submenu, move the cursor to the **List All Logical Unit Snapshots** option, and press Enter.

   The Cluster and Storage Pool Selection wizard starts.
3. Select the cluster name and press Enter.
4. Select the storage pool name and press Enter.
5. Press Enter to display all logical unit snapshots.

Rolling back to the logical unit snapshot

You can roll back to the logical unit snapshot in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu. Snapshots are images of a single logical unit or multiple logical units.

Before you begin

Note:

- If the logical unit is a rootvg device, you must shut down the client partition before you roll back to the logical unit snapshot.
- If the logical unit is a datavg device, stop access to all the volume groups in the virtual disk by using the **varyoffvg** command.

About this task

To roll back to a logical unit snapshot:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.

2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Roll Back to Snapshot** option, and press Enter.
3. Enter the cluster name, storage pool name, the snapshot to roll back to, and the list of logical units and press Enter.
4. Press Enter to roll back to the selected snapshot.
5. In the confirmation window that opens, press Enter to proceed with rolling back to the selected snapshot.

Deleting a logical unit snapshot

You can delete a logical unit snapshot in shared storage pools by using the Virtual I/O Server (VIOS) configuration menu. Snapshots are images of a single logical unit or multiple logical units.

About this task

To delete a logical unit snapshot:

Procedure

1. From the **Shared Storage Pools**, move the cursor to the **Manage Logical Units in Storage Pool** submenu and press Enter.
2. From the **Manage Logical Units in Storage Pool** submenu, move the cursor to the **Delete Snapshot** option and press Enter.
3. Enter the cluster name, storage pool name, the snapshot to delete, and the list of logical units. Press Enter.
4. Press Enter to delete the selected snapshot.
5. In the confirmation window that opens, press Enter to proceed with deleting the selected snapshot.

Getting started with Trusted Logging

Learn about using the Virtual I/O Server (VIOS) command line to configure the Trusted Logging capability for increased log security.

By using the PowerSC Trusted Logging capability, you can configure AIX logical partitions to write to log files that are stored on an attached VIOS. Data is transmitted to the VIOS directly through the hypervisor. Thus, configured network connectivity is not required between the client logical partitions and the VIOS on which the log files are stored.

The VIOS administrator can create and manage the log files by using the VIOS command-line interface. The following table lists the commands that can be used to configure and manage the Trusted Logging capability.

Table 39. Commands to configure and manage the Trusted Logging capability

Command	Description
chvlog	Changes the configuration of an existing virtual log.
chvlrepo	Changes the configuration of a virtual log repository.
lsvlog	Lists the currently defined virtual logs.
lsvlrepo	Lists the current configuration of the virtual log repositories.
mkvlog	Creates a new virtual log.
rmvlog	Removes an existing virtual log.

The Trusted Logging capability introduces the following concepts:

- Virtual log repositories

- Virtual logs
- Virtual log devices

These concepts are present in the VIOS as illustrated in the following figure. The virtual log devices are attached to virtual Small Computer Serial Interface (SCSI) adapters to expose the virtual log functions to client logical partitions. The virtual log devices are backed by virtual logs. Virtual logs are present in the VIOS file system as subdirectories within the virtual log repository. The virtual log repository is a directory in the VIOS file system.

The following figure shows the concepts of the Trusted Logging capability.

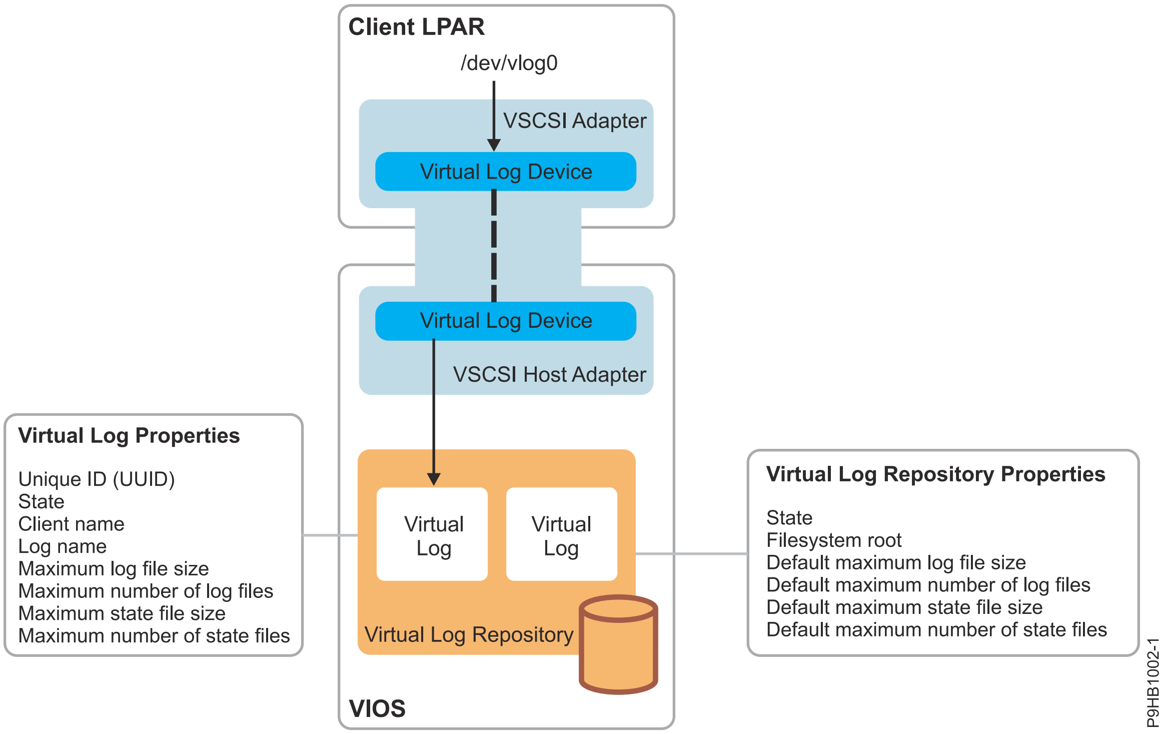

Related information

chvlog command

chvlrepo command

lsvlog command

lsvlrepo command

mkvlog command

rmvlog command

Virtual log repositories

Virtual log repositories are directories in the file system accessible by the Virtual I/O Server (VIOS). You can create one or more virtual logs in a virtual log repository.

Every VIOS has at least the local virtual log repository in the `/var/vio/vlogs` directory by default. If the VIOS is configured to use shared storage pools, there is another repository that is associated with each shared storage pool. When virtual logs are created, they are placed inside a specified virtual log repository. If an alternative repository is not specified, the local repository is used by default. The VIOS administrator can change the location of the local repository in the file system. However, shared storage pool repositories must reside in a fixed location.

Virtual logs

A virtual log is a directory in a virtual log repository.

The virtual log is used to store logs that are generated by an AIX logical partition. The properties of a virtual log can either be specified or inherited from the virtual log repository when the virtual log is created. The following table lists the virtual log properties.

Table 40. Virtual log properties

Property	Description
Unique ID (UUID)	Specifies the unique ID of the virtual log. This value is assigned when the virtual log is created and is retained permanently. If a logical partition is migrated to another system, the virtual log is re-created with the same configuration and unique ID on the destination Virtual I/O Server (VIOS) partition. For more information, see "Live Partition Mobility of virtual log devices" on page 174.
State	Indicates whether the virtual log can be attached to a client logical partition. It has the following possible values: • Enabled: Indicates that the virtual log can be attached to a client logical partition. • Migrated: Indicates that the virtual log is active on another VIOS after a migration operation. • Disabled: Indicates that the virtual log is not available for use by a client logical partition.
Client name	Indicates the name of the client. This property can be set to any value. However, typically all virtual logs that are intended for a particular client logical partition are assigned the same client name, for ease of administration. If a virtual log is created and attached to a client logical partition in a single operation, the VIOS attempts to obtain the host name from the client logical partition and use that as the client name if it is not specified on the command line.
Log name	Indicates the name of a virtual log. This property can be assigned any value by the administrator of the client logical partition, depending on the purpose, and must be provided when a new virtual log is created. For example, you can create two virtual logs, *audit* and *syslog*, for a given logical partition for the collection of audit and syslog data.
Maximum log file size	Specifies the maximum file size of the virtual log (in bytes).
Maximum number of log files	Specifies the maximum number of virtual log files.
Maximum state file size	Specifies the maximum size of the state file in bytes. A state file consists of additional information about when the virtual log devices were configured, opened, closed, and various other operations that might be of interest in an analysis of log activity.
Maximum number of state files	Specifies the maximum number of state files. A state file consists of additional information about when the virtual log devices were configured, opened, closed, and various other operations that might be of interest in an analysis of log activity.

Notes:

- The client name and log name properties also define the directory within the virtual log repository in which the log is stored. A virtual log repository contains a subdirectory for each client name. This subdirectory contains a directory for each log name. For example, with the local virtual log repository set to the default directory `/var/vio/vlogs`, a virtual log with the client name *lpar-01* and the log name *audit* stores the logs in the `/var/vio/vlogs/lpar-01/audit/` directory.

- If you rename the logical partition or change the host name, the client name property is not automatically updated. Use the **chvlog** command to change the value of the client name for the virtual log.

Each virtual log consists of the following types of information:

- Log data: Raw log data generated by the client logical partition. The log data is stored in files named in the *client_name_log_name.nnn* format.
- State data: Additional information about when the virtual log devices were configured, opened, closed, and various other operations that might be of interest in an analysis of log activity. This data is generated without any explicit user action. The state data is stored in files that are named in the *client_name_log_name.state.nnn* format.

In both cases, *nnn* starts at 000. The data is written to that file until the next write operation increases the size of the file to a value more than the maximum log file size. When the size of the file increases to a value more than the maximum log file size, *nnn* is incremented and a new file is created, overwriting any existing file with that name. Log data is written to the new file until *nnn* is incremented again and it reaches the limit specified in the properties of the virtual log. At this stage, *nnn* is reset to 000.

For example, consider a virtual log with the following properties:

```
Client name:                    lpar-01
Log name:                       audit
Maximum number of log files:    3
Maximum log file size:          2091216
Maximum number of state files:  2
Maximum state file size:        1000000
```

After a period of continued log generation, where the log files might have wrapped multiple times, the following directory contents are expected. The new log data is written to *lpar-01_audit.002* and the new state data are written to *lpar-01_audit.state.000*. For example, running `ls –l /var/vio/vlogs/lpar-01/audit` results in the following output:

```
-rw------- 1 root    system      2091216 May 25 18:28 lpar-01_audit.000
-rw------- 1 root    system      2091216 May 25 18:38 lpar-01_audit.001
-rw------- 1 root    system       752104 May 25 18:48 lpar-01_audit.002
-rw------- 1 root    system        16450 May 25 18:45 lpar-01_audit.state.000
-rw------- 1 root    system      1000000 May 21 07:23 lpar-01_audit.state.001
```

Virtual log devices

A virtual log device is a virtual target device on the Virtual I/O Server (VIOS), attached to a virtual Small Computer Serial Interface (SCSI) host adapter and backed by a virtual log.

By creating virtual log devices, virtual logs are made available to client logical partitions. The following sections describe the use of the local virtual log repositories.

See “Virtual log devices with shared storage pools” on page 174 topic for the commands that can also be used to work with virtual logs within a shared storage pool.

Configuring the virtual log repository

You can configure a virtual log repository by using the **chvlrepo** command. You can display the properties of the virtual log repositories by using the **lsvlrepo** command.

To configure or display the properties of a virtual log repository, use the following commands, as applicable:

- To display the current properties of virtual log repositories, enter the **lsvlrepo** command. Entering the **lsvlrepo -detail** command returns results similar to the following:

  ```
  Local Repository:
  State:                   enabled
  Repository Root:         /var/vio/vlogs
  Maximum Log Files:       10
  Maximum Log File Size:   2097152
  Maximum State Files:     10
  Maximum State File Size: 1048576
  ```

- To display this information in a custom format, use the **-field** flag. Specify a string with field names, which are separated by characters that are not alphanumeric to display a customized output. The output contains one line for every virtual log repository. For example, entering the `lsvlrepo -field "state-path lf"` command returns results similar to one of the following:

  -
    ```
    enabled-/tmp/vlogs/ 10
    ```

  -
    ```
    disabled-/var/vio/SSP/cTA1/D_E_F_A_U_L_T_061310/vlogs/ 3
    ```

  See lsvlrepo command for a list of all field names.

- To change the directory in which virtual logs are stored, enter the **chvlrepo** command. The virtual log repository directory cannot be changed if any virtual logs exist in the repository. To change the directory, enter the following command:

  ```
  chvlrepo -path /mnt/logs
  ```

- You can change properties, such as the default number and size of log files, by using other options of the **chvlrepo** command. See chvlrepo command for a list of all the options. For example, entering the following command changes the default values for virtual logs that are created in the local virtual log repository to have 4 log files, each up to 3 MB, and two state files, each up to 100 KB:

  ```
  chvlrepo -lf 4 -lfs 3M -sf 2 -sfs 100K
  ```

  Changing these default values does not change the configuration of existing virtual logs.

- You can also use the **chvlrepo** command to disable the repository to stop the creation of virtual logs. A virtual log repository cannot be disabled if there are any virtual logs in the repository. For example, entering the following command disables the repository:

  ```
  chvlrepo -state disabled
  ```

Creating a virtual log

You can create a virtual log and attach it to a virtual Small Computer Serial Interface (SCSI) host adapter by using the **mkvlog** command.

About this task

To create a virtual log and attach it to a virtual SCSI (VSCSI) host adapter, complete the following tasks:

Procedure

1. Enter the **mkvlog** command to create virtual logs. For example, entering the `mkvlog -name syslog -client lpar-01` command returns results similar to the following:

   ```
   Virtual log 00000000000000005b3f6b7cfcec4c67 created
   ```

   This command creates the *syslog* virtual log with the *lpar-01* client name and other properties that are inherited from the default values that are associated with the virtual log repository. The **mkvlog** command returns the UUID that has been assigned to the new virtual log.

2. Attach the virtual log that has been created to a VSCSI host adapter for use by a client logical partition. The VSCSI host adapter must not be configured to use the *Any Client Can Connect* mode. If you specify this mode, you cannot identify the logical partition that generated the log messages in the log files of the virtual log. For example, to attach the virtual log with UUID *00000000000000005b3f6b7cfcec4c67* to the VSCSI host adapter *vhost0*, enter the following command:

   ```
   mkvlog -uuid 00000000000000005b3f6b7cfcec4c67 -vadapter vhost0
   ```

   Results similar to the following are displayed:

   ```
   vtlog0 Available
   ```

Results

You can also create a virtual log and attach it to a VSCSI host adapter by using a single command instead of using the commands that are specified in step "1" on page 171 and "2" on page 172. For example, entering the `mkvlog -name audit -vadapter vhost1` command creates a new virtual log with the log name *audit*. This virtual log is attached to the VSCSI host adapter *vhost1*, with the client name set to the host name of the client logical partition that is attached to *vhost1*. Results similar to the following are displayed:

```
Virtual log 0000000000000000d96e956aa842d5f4 created
vtlog0 Available
```

Note: If the client logical partition is running, the client name does not need to be specified because the **mkvlog** command discovers the client name from the client logical partition.

Listing virtual logs or virtual log devices

You can list virtual logs or virtual log devices by using the **lsvlog** command.

To list virtual logs or virtual log devices, use the following commands, as applicable:

- To display the properties of virtual logs, enter the **lsvlog** command. For example, entering the **lsvlog** command returns results similar to the following:

  ```
  Client Name      Log Name         UUID                              VTD
  lpar-03          syslog           02392437473b6c552680a9ddd2fd8d06  vhost1/vtlog1
  lpar-02          syslog           956f8c1c25208091495c721e0796f456  vhost0/vtlog0
  lpar-01          audit            9705340b31a7883573a1cd04b2254efd
  lpar-01          syslog           b27a94a8e187ee5c917577c2a2df0268
  ```

- You can filter the output by using options such as **-uuid** to display only the log with a specific UUID. For example, entering the `lsvlog -uuid 02392437473b6c552680a9ddd2fd8d06` command returns results similar to the following:

  ```
  Client Name  Log Name  UUID                              VTD
  lpar-03      syslog    02392437473b6c552680a9ddd2fd8d06  vhost1/vtlog1
  ```

- To display all properties for each virtual log, use the **-detail** option. The virtual logs are displayed and are sorted by client name. For example, entering the `lsvlog -uuid 02392437473b6c552680a9ddd2fd8d06 -detail` command returns results similar to the following:

  ```
  Client Name: lpar-03
  Log Name:                     syslog
  UUID:                         02392437473b6c552680a9ddd2fd8d06
  Virtual Target Device:        vtlog1
  Parent Adapter:               vhost1
  State:                        enabled
  Logical Unit Address:         8100000000000000
  Log Directory:                /var/vio/vlogs/lpar-03/syslog
  Maximum Log Files:            10
  Maximum Log File Size:        1048576
  Maximum State Files:          10
  Maximum State File Size:      1048576
  ```

- To display this information in a custom format, use the **-field** option. Specify a string with field names that are separated by characters that are not alphanumeric. For example, entering the **`lsvlog -field "uuid\tsfs:sf"`** command lists all virtual logs. Results similar to the following are displayed:

  ```
  02392437473b6c552680a9ddd2fd8d06       1048576:10
  956f8c1c25208091495c721e0796f456       1048576:10
  9705340b31a7883573a1cd04b2254efd       1048576:5
  b27a94a8e187ee5c917577c2a2df0268       65536:20
  ```

Related information

lsvlog command

Reconfiguring virtual logs or virtual log devices

You can reconfigure virtual logs or virtual log devices by using the **`chvlog`** command.

To reconfigure virtual logs or virtual log devices, use the following commands, as applicable:

- To change the properties of a virtual log, enter the **`chvlog`** command. You can change the properties of virtual logs even if the virtual log is attached to a virtual log device on a virtual Small Computer Serial Interface (SCSI) adapter, and the changes are immediate.
- If the virtual log is connected to a virtual SCSI adapter, it can be specified by using the name of the virtual log device. For example, to change the log file size on the running virtual log device *vtlog0* to 2 MB, enter the `chvlog -dev vtlog0 -lfs 2M` command. Results similar to the following are displayed:

  ```
  Updated device.
  ```

- Regardless of whether a virtual log is connected to a virtual SCSI adapter, a virtual log can always be specified by using the UUID of the virtual log. For example, to change the state of the virtual log with UUID *00000000000000003cee6408c885d677* to disabled, enter the `chvlog -uuid 00000000000000003cee6408c885d677 -state disabled` command. Results similar to the following are displayed.

  ```
  Updated device.
  ```

- The state property for a virtual log controls whether the virtual log can be connected to a virtual SCSI adapter. Therefore, it is not valid to change the state property when the virtual log is attached to a virtual log device. For example, to change the state of the virtual log with UUID *00000000000000003cee6408c885d677* to *disabled* when it is connected to a virtual SCSI host adapter, enter the `chvlog -uuid 00000000000000003cee6408c885d677 -state disabled` command. Results similar to the following are displayed:

  ```
  To change the state, the virtual log must not be connected to a device.
  ```

  If you enter the **`lsvlog`** command, the VTD column is blank for this virtual log.

  Note: To delete the virtual log device while retaining the virtual log, use the **`rmvlog -d`** command.

Removing virtual logs or virtual log devices

You can use the **`rmvlog`** command to remove virtual logs or virtual log devices from a virtual Small Computer Serial Interface (SCSI) adapter, or to unconfigure a virtual log device. The virtual log can be specified by using the UUID or by the associated virtual log device name, if it exists.

To remove virtual log devices or virtual logs, use the following commands, as applicable:

- To change the specified virtual log device from the *Available* state to the *Defined* state, enter the **`rmvlog`** command. To specify the virtual log device by name, use the **-dev** option. For example, entering `rmvlog -dev vtlog0` returns results similar to the following:

  ```
  vtlog0 Defined
  ```

- To specify the virtual log device, use the **-uuid** option. When you use this option, the virtual log device that is associated with a virtual log and the specified UUID is changed. For example, entering `rmvlog -uuid 0000000000000000a3e4dd0ba75972c2` returns results similar to the following:

  ```
  vtlog0 Defined
  ```

- To remove the specified virtual log device, specify the **-d** option in addition to either the **-dev** or **-uuid** option. When you use the **-d** option, the virtual log device is deleted. However, the virtual log and all associated properties and data are retained. For example, entering the `rmvlog -dev vtlog0 -d` returns results similar to the following:

  ```
  vtlog0 deleted
  ```

- To remove the virtual log device and the virtual log, specify the **-db** option. When you use this option, the data is still retained. For example, entering the `rmvlog -uuid 9705340b31a7883573a1cd04b2254efd -db` command returns results similar to the following:

  ```
  Virtual log 9705340b31a7883573a1cd04b2254efd deleted.
  ```

- To remove the virtual log device, the virtual log, and any log files that are associated with the virtual log, specify the **-dbdata** option. For example, entering the `rmvlog -dev vtlog0 -dbdata` command returns results similar to the following:

  ```
  vtlog1 deleted
  Virtual log 02392437473b6c552680a9ddd2fd8d06 deleted.
  Log files deleted.
  ```

Live Partition Mobility of virtual log devices

When a client logical partition is moved from one host system to another during Live Partition Mobility, new virtual log devices are created on the destination Virtual I/O Server (VIOS).

When you do not use Shared Storage Pools, these new virtual logs are independent of the virtual logs on the source VIOS. The configuration data of the source virtual log without the log file content is copied to the destination virtual log during migration. After migration, the source virtual log is placed in the migrated state to indicate that the virtual log is no longer active on the system and that it has been moved to another system. If you use a migration operation to move the client logical partition back to the original host system, and you select the original VIOS to host the virtual logs of the logical partition, the existing virtual log is moved back to the enabled state.

Virtual log devices with shared storage pools

You can use the Trusted Logging feature to direct log data to a file system shared across Virtual I/O Server (VIOS) logical partitions.

By using the Trusted Logging feature with shared storage pools, you can obtain a single view of logical partition activity across several separate systems.

Benefits of using virtual log devices with shared storage pools

Using virtual log devices with shared storage pools provides multipath logs on a single system and Live Partition Mobility of virtual logs.

You can use the trusted log feature to direct log data to a file system shared across more than one Virtual I/O Server (VIOS) and obtain a single view of logical partition activity across several separate systems. This feature provides the following benefits:

- Multipath logs on a single system: By using virtual logs in shared storage pools, more than one VIOS on a single host can make the same virtual log available to a client logical partition through different virtual Small Computer Serial Interface (SCSI) host adapters. The client logical partition detects the multipath arrangement and tolerates the deactivation of a single VIOS by failing over to an alternative path, without losing log data.

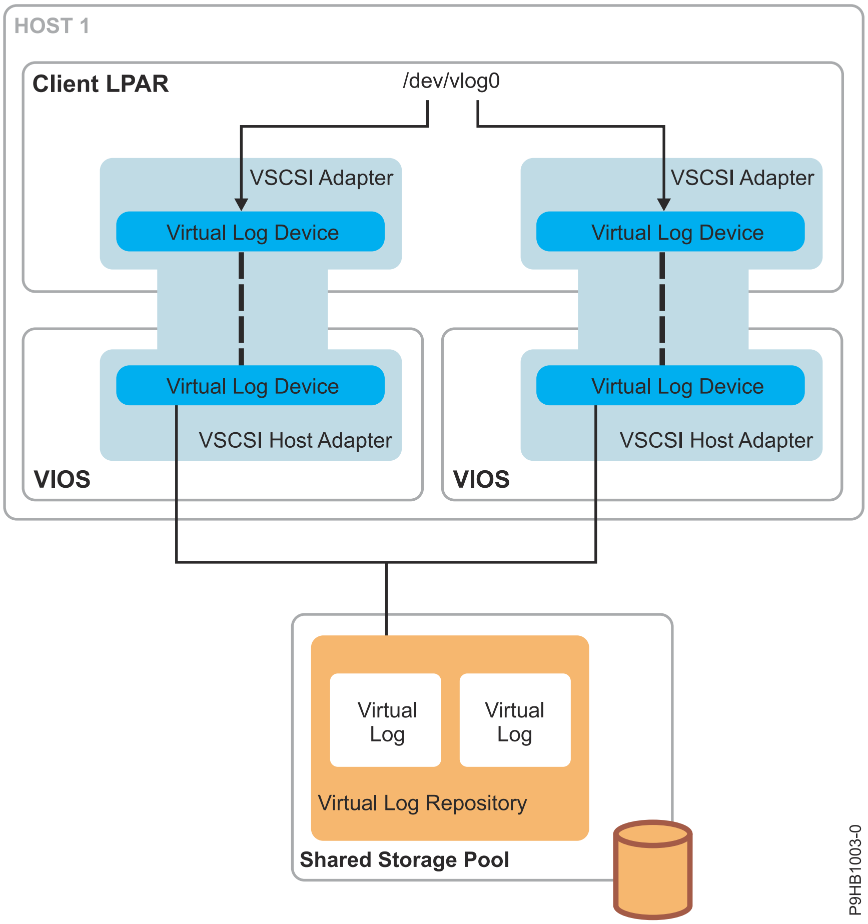

- Live Partition Mobility of virtual logs: When VIOS logical partitions on two different hosts have visibility of the same shared storage pool virtual log repository, a migration operation is able to continuously write to a single set of log files inside the shared storage pool, rather than to two different local virtual log repositories. Thus, in contrast to Live Partition Mobility with local virtual log repositories where the log files are split across two file systems, a single log file continues to be written across the migration operation.

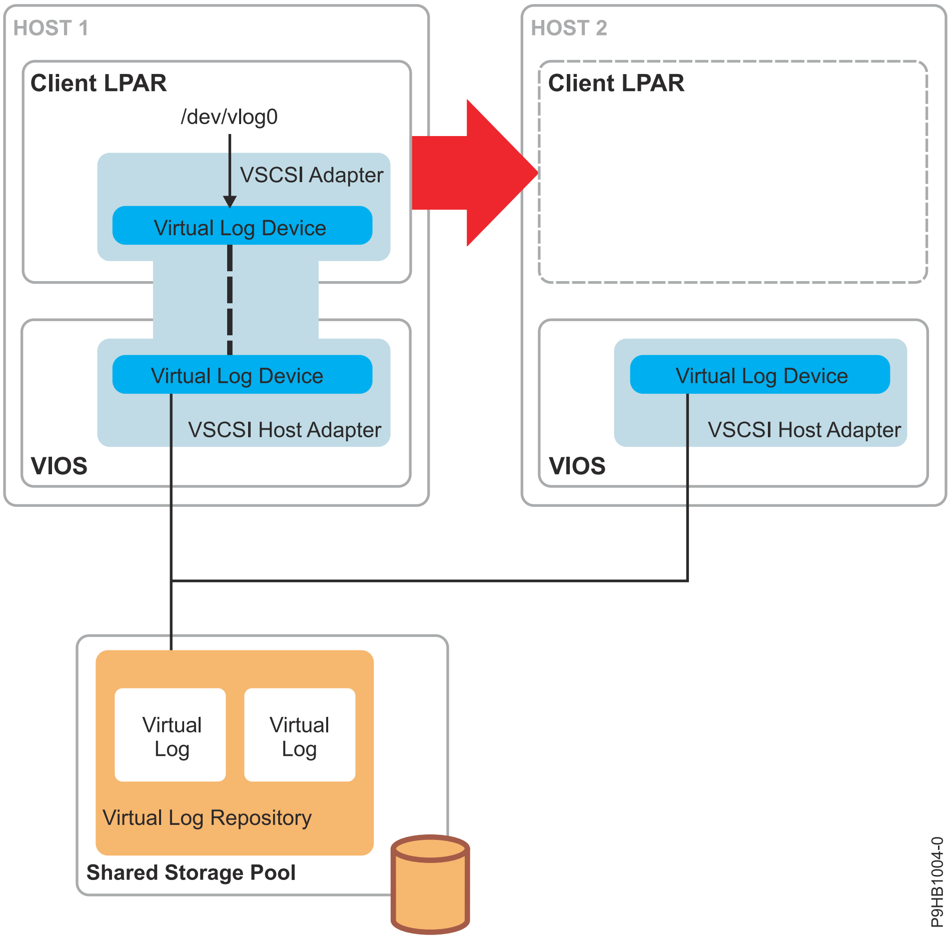

Using virtual log devices with shared storage pools

Learn about using virtual log devices with shared storage pools.

About this task

To use virtual logs with shared storage pools, the VIOS logical partitions must be clustered together. For instructions, see "Configuring the system to create shared storage pools" on page 122. This process creates a shared storage pool, the name of which is used in virtual log commands to operate on virtual logs within that shared storage pool. To create a virtual log inside a shared storage pool, complete the following tasks:

Procedure

1. Run the **mkvlog** command as described in "Creating a virtual log" on page 171. In addition, specify the **-sp** option to indicate the shared storage pool to use.

   For example, entering the `mkvlog -sp spool1 -name syslog -client lpar-01` command returns results similar to the following:

   ```
   Virtual log f5dee41bf54660c2841c989811de41dd created
   ```

2. Attach the virtual log that was created in the shared storage pool to virtual Small Computer Serial Interface (SCSI) adapters. For example, entering the `mkvlog -uuid f5dee41bf54660c2841c989811de41dd -vadapter vhost0` command returns results similar to the following:

   ```
   vtlog1 Available
   ```

Results

Notes:

- The **lsvlog**, **chvlog**, and **rmvlog** commands operate on virtual logs in shared storage pools in the same way that they operate on virtual logs in the local virtual log repository. However, the **chvlog** command cannot be used to change virtual logs that are currently connected to virtual log devices anywhere in the cluster. The virtual log devices must be removed before changes can be made to the virtual log configuration.
- Additionally, the root path to a shared storage pool virtual log repository cannot be changed. The location is decided by the mount point of the shared storage pool on the Virtual I/O Server (VIOS).

Each shared storage pool has a separate virtual log repository with a separate set of default properties that are inherited by virtual logs that are created within that virtual log repository. By default, the **lsvlrepo** command displays the properties of all virtual log repositories. You can use the **-local** and **-sp** options to display the properties of a specific virtual log repository.

Getting started with Trusted Firewall

Learn about using the Trusted Firewall feature that is supported on the PowerSC Editions. You can use this feature to perform intervirtual LAN routing functions by using the Security Virtual Machine (SVM) kernel extension.

With Virtual I/O Server (VIOS) Version 2.2.1.4, or later, you can configure and manage the Trusted Firewall feature. By using this feature, logical partitions on different VLANs of the same server can communicate through the shared Ethernet adapter. The shared Ethernet adapter calls the intervirtual LAN routing functions through the SVM kernel extension.

The SVM kernel extension consists of the following intervirtual LAN routing functions:

- Layer 3 routing: VLANs represent different logical networks. Therefore, a layer 3 router is required to connect the VLANs.
- Network filtering rules: Network filtering rules are required to permit, deny, or route intervirtual LAN network traffic. Network filtering rules can be configured by using the VIOS command-line interface.

The following table lists the commands that can be used to configure and manage the Trusted Firewall feature by using the VIOS command-line interface.

Table 41. Commands to configure and manage the Trusted Firewall feature

Command	Description
chvfilt	Changes the definition of a VLAN-crossing filter rule in the filter rule table.
genvfilt	Adds a filter rule for the VLAN-crossing between logical partitions on the same Power Systems server.
lsvfilt	Lists the VLAN-crossing filter rules and their status.
mkvfilt	Activates the VLAN-crossing filter rules that are defined by the **genvfilt** command.
rmvfilt	Removes the VLAN-crossing filter rules from the filter table.
vlantfw	Displays or clears the IP and Media Access Control (MAC) mappings.

Related reference

PowerSC

Trusted Firewall

Related information

chvfilt command

genvfilt command

lsvfilt command

mkvfilt command

rmvfilt command

vlantfw command

Configuring virtual Ethernet on the Virtual I/O Server

You can configure virtual Ethernet devices by deploying a system plan, create and configure a Shared Ethernet Adapter (SEA), and configure a Link Aggregation device.

About this task

For better performance, you can configure the IP address by using the SEA directly as follows:

- If the VLAN is the same as the PVID, you can configure the IP address by using the SEA interface.
- If the VLAN is not the same as the PVID, you can create a VLAN pseudo-device with the VLAN ID and assign the IP address by using the interface of the pseudo-device.

However, if the SEA fails, the IP address that is configured on it would go down, as a result.

In SEA failover configuration for higher availability, you can create a virtual adapter with the Port VLAN ID (PVID) of the corresponding Virtual LAN (VLAN) and configure the IP address by using the interface of that virtual adapter.

Creating a virtual Ethernet adapter with the HMC Version 7 graphical interface

Using the Hardware Management Console (HMC), Version 7 Release 3.4.2 or later, you can create a virtual Ethernet adapter on a Virtual I/O Server (VIOS). With a virtual Ethernet adapter, client logical partitions can access the external network without having to own a physical Ethernet adapter.

Before you begin

If you plan to use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet), ensure that the Logical Host Ethernet adapter (LHEA) on the Virtual I/O Server is set to promiscuous mode.

For more information about adding a virtual network and creating virtual Ethernet adapters when the HMC is at version 8.7.0, or later, see The Add Virtual Network wizard.

Note: For HMC versions before Version 7, Release 3.4.2, you must use the VIOS command-line interface to configure the adapter.

What to do next

When you have completed the steps, configure the Shared Ethernet Adapter with the Virtual I/O Server command-line interface, or the Hardware Management Console graphical interface, Version 7 Release 3.4.2, or later.

Related concepts

Setting the SR-IOV Ethernet Logical Port to promiscuous mode

To use a Shared Ethernet Adapter with an SR-IOV Ethernet Logical Port, you must set the SR-IOV Ethernet Logical Port to have the promiscuous permission. You select the promiscuous permission for an SR-IOV Logical port, when you assign an SR-IOV Logical Port to a logical partition or logical partition profile, or when you add an SR-IOV Logical Port to a logical partition dynamically.

Related tasks

Configuring a Shared Ethernet Adapter with the Virtual I/O Server command-line interface

To configure a shared Ethernet adapter (SEA) with Hardware Management Console versions before 7, Release 3.4.2, you must use the Virtual I/O Server command-line interface.

Setting the LHEA to promiscuous mode

To use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet), you must set the Logical Host Ethernet Adapter (LHEA) to promiscuous mode.

Configuring a Shared Ethernet Adapter with the Virtual I/O Server command-line interface
To configure a shared Ethernet adapter (SEA) with Hardware Management Console versions before 7, Release 3.4.2, you must use the Virtual I/O Server command-line interface.

Setting the SR-IOV Ethernet Logical Port to promiscuous mode

To use a Shared Ethernet Adapter with an SR-IOV Ethernet Logical Port, you must set the SR-IOV Ethernet Logical Port to have the promiscuous permission. You select the promiscuous permission for an SR-IOV Logical port, when you assign an SR-IOV Logical Port to a logical partition or logical partition profile, or when you add an SR-IOV Logical Port to a logical partition dynamically.

To assign an SR-IOV logical port, complete the following steps:

1. In the Create LPAR wizard page, click **SR-IOV Logical Ports**.
2. Click **Actions** > **Create Logical Port** > **Ethernet Logical Port**.
3. In the Add Ethernet Logical Port page, select the physical port for the logical port.
4. Click **OK**.
5. Click the **General** tab of the Logical Port Properties page.
    a. In the permissions area of the **General** tab, enable the Promiscuous options, by selecting the appropriate check box.

For more information about dynamically adding SR-IOV logical ports when the HMC is at version 8.7.0, or later, see Adding SR-IOV logical ports.

Setting the LHEA to promiscuous mode

To use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet), you must set the Logical Host Ethernet Adapter (LHEA) to promiscuous mode.

Before you begin

Before you start, use the Hardware Management Console (HMC) to determine the physical port of the Host Ethernet Adapter that is associated with the Logical Host Ethernet port. Determine this information for the Logical Host Ethernet port that is the real adapter of the Shared Ethernet Adapter on the Virtual I/O Server. You can find this information in the partition properties of the Virtual I/O Server, and the managed system properties of the server on which the Virtual I/O Server is located.

About this task

For more information about adding SR-IOV logical ports and setting the Logical Host Ethernet port (that is the real adapter of the Shared Ethernet Adapter) to promiscuous mode, when the HMC is at version 8.7.0, or later, see Adding SR-IOV logical ports.

Configuring a Shared Ethernet Adapter with the Virtual I/O Server command-line interface

To configure a shared Ethernet adapter (SEA) with Hardware Management Console versions before 7, Release 3.4.2, you must use the Virtual I/O Server command-line interface.

Before you begin

In SEA, quality of service (QoS) is provided per SEA thread. By default, SEA runs in thread mode with seven threads. When SEA receives traffic, it routes the traffic to a thread, based on source and destination information. If the QoS mode is enabled, each thread further queues the traffic, based on the VLAN tag priority, to the appropriate priority queue associated with the selected thread. Queued traffic for a particular thread is serviced in the order of higher to lower priority. All threads handle all priorities.

Note: SEA QoS does not assure bandwidth for a particular priority. The packets are prioritized by each thread locally, not across the multiple SEA threads globally.

The SEA QoS is effective when all SEA threads are handling traffic, such that when an SEA thread is scheduled to run, it services higher priority traffic before servicing the lower priority traffic. An SEA QoS is not effective when the higher and lower priority traffic is spread across different threads.

Before you can configure an SEA, you must first create the virtual Ethernet trunk adapter by using the Hardware Management Console (HMC).

About this task

You can configure an SEA with the Virtual I/O Server command-line interface.

Procedure

1. Verify that the virtual Ethernet trunk adapter is available by running the following command:

   ```
   lsdev -virtual
   ```

2. Identify the appropriate physical Ethernet adapter that is used to create the SEA by running the following command:

   ```
   lsdev -type adapter
   ```

   Notes:

   - Ensure that TCP/IP is not configured on the interface for the physical Ethernet adapter. If TCP/IP is configured, the **mkvdev** command in the next step fails.
   - You can also use a Link Aggregation, or Etherchannel, device as the SEA.
   - If you plan to use the Host Ethernet Adapter or Integrated Virtual Ethernet with the SEA, ensure that you use the Logical Host Ethernet adapter to create the SEA.

3. Configure an SEA by running the following command:

   ```
   mkvdev -sea target_device -vadapter virtual_ethernet_adapters \
   -default DefaultVirtualEthernetAdapter -defaultid SEADefaultPVID
   ```

   Where:

   DefaultVirtualEthernetAdapter
      The default virtual Ethernet adapter used to handle untagged packets. If you have only one virtual Ethernet adapter for this logical partition, use it as the default.

   SEADefaultPVID
      The PVID associated with your default virtual Ethernet adapter.

   target_device
      The physical adapter that is being used as part of the SEA device.

   virtual_ethernet_adapters
      The comma-separated list of the virtual Ethernet adapters that are used as a part of the SEA device.

   For example:

   - To create an SEA`ent3` with `ent0` as the physical Ethernet adapter (or Link Aggregation) and `ent2` as the only virtual Ethernet adapter (defined with a PVID of `1`), type the following command:

     ```
     mkvdev -sea ent0 -vadapter ent2 -default ent2 -defaultid 1
     ```

   - To obtain the value for the `SEADefaultPVID` attribute in the **mkvdev** command, type the following command:

     ```
     entstat -all ent2 | grep "Port VLAN ID:"
     ```

Output similar to the following example is displayed:

```
Port VLAN ID: 1
```

4. Verify that the SEA was created by running the following command:

```
lsdev -virtual
```

5. Do you plan to access the Virtual I/O Server from the network with the physical device used to create the SEA?

   - Yes: Go to step “6” on page 181.
   - No: You are finished with this procedure and can skip the remaining steps.

6. Do you plan to set bandwidth apportioning by defining a quality of service (QoS)?

   - Yes: Go to step 11 to enable the SEA device to prioritize traffic.
   - No: Go to step 9 to configure a TCP/IP connection.

7. Do you plan to define IP addresses on any VLANs other than the VLAN specified by the PVID of the SEA?

   - Yes: Go to step “8” on page 181 to create VLAN pseudo-devices.
   - No: Go to step “9” on page 181 to configure a TCP/IP connection.

8. To configure VLAN pseudo-devices, complete the following steps:

   a) Create a VLAN pseudo-device on the SEA by running the following command:

   ```
   mkvdev -vlan TargetAdapter -tagid TagID
   ```

   Where:

   - *TargetAdapter* is the SEA.
   - *TagID* is the VLAN ID that you defined when you created the virtual Ethernet adapter associated with the SEA.

   For example, to create a VLAN pseudo-device by using the SEA `ent3` that you created with a VLAN ID of 1, type the following command:

   ```
   mkvdev -vlan ent3 -tagid 1
   ```

   b) Verify that the VLAN pseudo-device was created by running the following command:

   ```
   lsdev -virtual
   ```

   c) Repeat this step for any additional VLAN pseudo-devices that you need.

9. Run the following command to configure the first TCP/IP connection.

   The first connection must be on the same VLAN and logical subnet as the default gateway.

   ```
   mktcpip -hostname Hostname -inetaddr Address -interface Interface -netmask \
   SubnetMask -gateway Gateway -nsrvaddr NameServerAddress -nsrvdomain Domain
   ```

   Where:

   - *Hostname* is the host name of the Virtual I/O Server
   - *Address* is the IP address that you want to use for the TCP/IP connection
   - *Interface* is the interface that is associated with either the SEA device or a VLAN pseudo-device. For example, if the SEA device is ent3, the associated interface is en3.
   - *Subnetmask* is the subnet mask address for your subnet.
   - *Gateway* is the gateway address for your subnet.
   - *NameServerAddress* is the address of your domain name server.
   - *Domain* is the name of your domain.

If you do not have more VLANs, then you are finished with this procedure and can skip the remaining steps.

10. Run the following command to configure more TCP/IP connections:

```
chdev -dev interface -perm -attr netaddr=IPaddress netmask=netmask
state=up
```

While using this command, enter the interface (en*X*) associated with either the SEA device or the VLAN pseudo-device.

11. Enable the SEA device to prioritize traffic. Client logical partitions must insert a VLAN priority value in their VLAN header. For AIX clients, a VLAN pseudo-device must be created over the Virtual I/O Ethernet adapter, and the VLAN priority attribute must be set (the default value is 0). Do the following steps to enable traffic prioritization on an AIX client: For AIX clients, a VLAN pseudo-device must be created over the Virtual I/O Ethernet adapter, and the VLAN priority attribute must be set (the default value is 0). Do the following steps to enable traffic prioritization on an AIX client:

Note:

- While configuring QoS on the VLAN devices, you can also configure the QoS priority for a virtual Ethernet adapter by using the Hardware Management Console.
- You can also configure VLANs on Linux logical partitions. For more information, see the documentation for the Linux operating system.

a) Set the SEA qos_mode attribute to either strict or loose mode. Use one of the following commands: `chdev -dev <SEA device name> -attr qos_mode=strict` or `chdev -dev <SEA device name> -attr qos_mode=loose`.

For more information about the modes, see SEA.

b) From the HMC, create a Virtual I/O Ethernet Adapter for the AIX client with all of the tagged VLANs that are required (specified in the `Additional VLAN ID` list).

Packets that are sent over the default VLAN ID (specified in the **Adapter ID** or **Virtual LAN ID** field) are not tagged as VLAN; therefore, a VLAN priority value cannot be assigned to them.

c) On the AIX client, run the **`smitty vlan`** command.

d) Select **Add a VLAN**.

e) Select the name of the Virtual I/O Ethernet Adapter created in step 1.

f) In the VLAN Tag ID attribute, specify one of the tagged VLANs that are configured on the Virtual I/O Ethernet adapter that you created in step 1.

g) Specify an attribute value (0 - 7) in the VLAN Priority attribute, which corresponds to the importance the VIOS gives to the traffic sent over that VLAN pseudo-device.

h) Configure the interface over the VLAN pseudo-device that is created in step 6.

a) Set the SEA qos_mode attribute to either strict or loose mode. Use one of the following commands: `chdev -dev <SEA device name> -attr qos_mode=strict` or `chdev -dev <SEA device name> -attr qos_mode=loose`.

For more information about the modes, see SEA.

b) From the HMC, create a Virtual I/O Ethernet Adapter for the AIX client with all of the tagged VLANs that are required (specified in the `Additional VLAN ID` list).

Packets sent over the default VLAN ID (specified in the **Adapter ID** or **Virtual LAN ID** field) will not be tagged as VLAN; therefore, a VLAN priority value cannot be assigned to them.

c) On the AIX client, run the **`smitty vlan`** command.

d) Select **Add a VLAN**.

e) Select the name of the Virtual I/O Ethernet Adapter created in step 1.

f) In the VLAN Tag ID attribute, specify one of the tagged VLANs that are configured on the Virtual I/O Ethernet Adapter that you created in step 1.

g) Specify an attribute value (0 - 7) in the VLAN Priority attribute, which corresponds to the importance the VIOS must give to the traffic sent over that VLAN pseudo-device.

h) Configure the interface over the VLAN pseudo-device that is created in step 6.

Traffic sent over the interface created in step 7 is tagged as VLAN and its VLAN header has the VLAN priority value that is specified in step 6. When this traffic is bridged by a SEA that has been enabled for bandwidth apportioning, the VLAN priority value is used to determine how quickly it must be sent in relation to other packets at different priorities.

Results

The Shared Ethernet Adapter is now configured. After you configure the TCP/IP connections for the virtual adapters on the client logical partitions by using the client logical partitions' operating systems, those logical partitions can communicate with the external network.

Related concepts

Shared Ethernet Adapter failover
Shared Ethernet Adapter failover provides redundancy by configuring a backup Shared Ethernet Adapter on a different Virtual I/O Server logical partition that can be used if the primary Shared Ethernet Adapter fails. The network connectivity in the client logical partitions continues without disruption.

Shared Ethernet Adapters
With Shared Ethernet Adapters on the Virtual I/O Server logical partition, virtual Ethernet adapters on client logical partitions can send and receive outside network traffic.

Related information

Creating a shared Ethernet adapter for a VIOS logical partition using the HMC

Virtual I/O Server commands

Creating a virtual Ethernet adapter by using HMC Version 7

Creating a shared Ethernet adapter for a Virtual I/O Server logical partition using the HMC Version 7, release 3.4.2 or later

Configuring a Link Aggregation or Etherchannel device

Configure a Link Aggregation device, also called an Etherchannel device, by using the **mkvdev** command. A Link Aggregation device can be used as the physical Ethernet adapter in the Shared Ethernet Adapter configuration.

About this task

Configure a Link Aggregation device by typing the following command:

```
mkvdev -lnagg TargetAdapter ... [-attr Attribute=Value ...]
```

For example, to create Link Aggregation device ent5 with physical Ethernet adapters ent3, ent4, and backup adapter ent2, type the following:

```
mkvdev -lnagg ent3,ent4 -attr backup_adapter=ent2
```

After the Link Aggregation device is configured, you can add adapters to it, remove adapters from it, or modify its attributes by using the **cfglnagg** command.

Assigning the virtual Fibre Channel adapter to a physical Fibre Channel adapter

To enable N-Port ID Virtualization (NPIV) on managed systems, connect the virtual Fibre Channel adapter on the Virtual I/O Server logical partition to a physical port on a physical Fibre Channel adapter.

Before you begin

Before you start, verify that the following statements are true:

- Verify that you have created the virtual Fibre Channel adapters on the Virtual I/O Server logical partition and associated them with virtual Fibre Channel adapters on the client logical partition.
- Verify that you have created the virtual Fibre Channel adapters on each client logical partition and associated them with a virtual Fibre Channel adapter on the Virtual I/O Server logical partition.

About this task

After the virtual Fibre Channel adapters are created, you need to connect the virtual Fibre Channel adapter on the Virtual I/O Server logical partition to the physical ports of the physical Fibre Channel adapter. The physical Fibre Channel adapter must be connected to the physical storage that you want the associated client logical partition to access.

Tip: If you are using the HMC, Version 7 Release 3.4.2 or later, you can use the HMC graphical interface to assign the virtual Fibre Channel adapter on a Virtual I/O Server to a physical Fibre Channel adapter.

To assign the virtual Fibre Channel adapter to a physical port on a physical Fibre Channel adapter, complete the following steps from the Virtual I/O Server command-line interface:

Procedure

1. Use the **`lsnports`** command to display information for the available number of NPIV ports and available worldwide port names (WWPNs).

   For example, running **`lsnports`** returns results similar to the following:

   ```
   Name      Physloc                    fabric    tports    aports    swwpns    awwpns
   -----------------------------------------------------------------------------------
   fcs0     U789D.001.DQDMLWV-P1-C1-T1  1         64        64         2048      2047
   fcs1     U787A.001.DPM0WVZ-P1-C1-T2  1         63        62         504       496
   ```

   Note: If there are no NPIV ports in the Virtual I/O Server logical partition, the error code `E_NO_NPIV_PORTS(62)` is displayed.
2. To connect the virtual Fibre Channel adapter on the Virtual I/O Server logical partition to a physical port on a physical Fibre Channel adapter, run the **vfcmap** command: `vfcmap -vadapter` *`virtual Fibre Channel adapter`* `-fcp` *`Fibre Channel port name`*

   where,

   - *Virtual Fibre Channel adapter* is the name of the virtual Fibre Channel adapter that is created on the Virtual I/O Server logical partition.
   - *Fibre Channel port name* is the name of the physical Fibre Channel port.

     Note: If no parameter is specified with the **-fcp** flag, the command unmaps the virtual Fibre Channel adapter from the physical Fibre Channel port.
3. Use the **`lsmap`** command to display the mapping between virtual host adapters and the physical devices to which they are backed. To list NPIV mapping information, type: `lsmap -all -npiv`.

   The system displays a message similar to the following:

   ```
   Name      Physloc                    ClntID    ClntName  ClntOS
   ---------------------------------------------------------------
   vfchost0  U8203.E4A.HV40026-V1-C12   1         HV-40026   LinuxAIXAIX

   Status:NOT_LOGGED_IN
   FC name:fcs0                     FC loc code:U789C.001.0607088-P1-C5-T1
   Ports logged in:0
   Flags:1 <not_mapped, not_connected>
   VFC client name:             VFC client DRC:
   ```

What to do next

When you are finished, consider the following tasks:

- For each logical partition, verify that both WWPNs are assigned to the same physical storage and have the same level of access on the storage area network (SAN). For instructions, see the IBM System Storage® SAN Volume Controller.

  Note: To determine the WWPNs that are assigned to a logical partition, use the Hardware Management Console (HMC) to view the partition properties or partition profile properties of the client logical partition.
- If you later need to remove the connection between the virtual Fibre Channel adapter that is created on the Virtual I/O Server logical partition and the physical port, you can do so by using the **vfcmap** command and not specifying a parameter for the **-fcp** flag.

Related information

Configuring a virtual Fibre Channel adapter

Changing virtual Fibre Channel by using the Hardware Management Console

Virtual I/O Server commands

Configuring the IBM Tivoli agents and clients on the Virtual I/O Server

You can configure and start the IBM Tivoli Monitoring agent, IBM Tivoli Usage and Accounting Manager, the IBM Tivoli Storage Manager client, and the Tivoli Storage Productivity Center agents.

Related concepts

IBM Tivoli software and the Virtual I/O Server
Learn about integrating the Virtual I/O Server into your Tivoli environment for IBM Tivoli Application Dependency Discovery Manager, IBM Tivoli Monitoring, IBM Tivoli Storage Manager, IBM Tivoli Usage and Accounting Manager, IBM Tivoli Identity Manager, and Tivoli Storage Productivity Center.

Related information

cfgsvc command

Configuring the IBM Tivoli Monitoring agent

You can configure and start the IBM Tivoli Monitoring agent on the Virtual I/O Server.

Before you begin

With Tivoli Monitoring System Edition for IBM Power Systems, you can monitor the health and availability of multiple Power Systems servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. IBM Tivoli Monitoring System Edition for Power Systems gathers data from the Virtual I/O Server, including data about physical volumes, logical volumes, storage pools, storage mappings, network mappings, real memory, processor resources, mounted file system sizes, and so on. From the Tivoli Enterprise Portal, you can view a graphical representation of the data, use predefined thresholds to alert you on key metrics, and resolve issues based on suggestions that are provided by the Expert Advice feature of Tivoli Monitoring.

Before you start, complete the following tasks:

- Ensure that the Virtual I/O Server is running fix pack 8.1.0. For instructions, see "Updating the Virtual I/O Server" on page 205.
- Verify that you are a super administrator of the HMC.
- Verify that you are the prime administrator of the Virtual I/O Server.

About this task

To configure and start the monitoring agent, complete the following steps:

Procedure

1. List all of the available monitoring agents by using the **lssvc** command.

   For example,

```
$lssvc
ITM_premium
```

2. Based on the output of the **lssvc** command, decide which monitoring agent you want to configure. For example, `ITM_premium`
3. List all of the attributes that are associated with the monitoring agent by using the **cfgsvc** command. For example:

   ```
   $cfgsvc –ls ITM_premium
    HOSTNAME
   RESTART_ON_REBOOT
   MANAGING_SYSTEM
   ```

4. Configure the monitoring agent with its associated attributes by using the **cfgsvc** command:

   ```
   cfgsvc ITM_agent_name -attr Restart_On_Reboot=value hostname=name_or_address1
   managing_system=name_or_address2
   ```

   Where:

   - *ITM_agent_name* is the name of the monitoring agent. For example, ITM_premium.
   - *value* must be either TRUE of FALSE as follows:
     - TRUE: *ITM_agent_name* restarts whenever the Virtual I/O Server restarts
     - FALSE: *ITM_agent_name* does not restart whenever the Virtual I/O Server restarts
   - *name_or_address1* is either the host name or IP address of the Tivoli Enterprise Monitoring Server server to which *ITM_agent_name* sends data.
   - *name_or_address2* is either the host name of IP address of the Hardware Management Console (HMC) attached to the managed system on which the Virtual I/O Server with the monitoring agent is located.

   For example:

   ```
   cfgsvc ITM_premium –attr Restart_On_Reboot=TRUE hostname=tems_server
   managing_system=hmc_console
   ```

   In this example, the ITM_premium monitoring agent is configured to send data to tems_server, and to restart whenever the Virtual I/O Server restarts.
5. Start the monitoring agent by using the **startsvc** command. For example:

   ```
   startsvc ITM_premium
   ```

6. From the HMC, complete the following steps so that the monitoring agent can gather information from the HMC.

   Note: After you configure a Secure Shell connection for one monitoring agent, you do not need to configure it again for any additional agents.

   a) Determine the name of the managed system on which the Virtual I/O Server with the monitoring agent is located.

   b) Obtain the public key for the Virtual I/O Server by running the following command:

   ```
   viosvrcmd -m managed_system_name -p vios_name -c "cfgsvc -key ITM_agent_name"
   ```

   Where:

   - *managed_system_name* is the name of the managed system on which the Virtual I/O Server with the monitoring agent or client is located.
   - *vios_name* is the name of the Virtual I/O Server logical partition (with the monitoring agent) as defined on the HMC.
   - *ITM_agent_name* is the name of the monitoring agent. For example, `ITM_premium`.

c) Update the authorized_key2 file on the HMC by running the **mkauthkeys** command:

```
mkauthkeys --add public_key
```

where, *public_key* is the output from the **viosvrcmd** command in step 6b.

For example:

```
$ viosvrcmd -m commo126041 -p VIOS7 -c "cfgsvc ITM_premium -key"
ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAvjDZ
  sS0guWzfzfp9BbweG0QMXv1tbDrtyWsgPbA2ExHA+xduWA51K0oFGarK2F
  C7e7NjKW+UmgQbrh/KSyKKwozjp4xWGNGhLmfan85ZpFR7wy9UQG1bLgXZ
  xYrY7yyQQQODjvwosWAfzkjpG3iW/xmWD5PKLBmob2QkKJbxjne+wqGwHT
  RYDGIiyhCBIdfFaLZgkXTZ2diZ98rL8LIv3qb+TsM1B28AL4t+1OGGeW24
  2lsB+8p4kamPJCYfKePHo67yP4NyKyPBFHY3TpTrca4/y1KEBT0Va3Pebr
  5JEIUvWYs6/RW+bUQk1Sb6eYbcRJFHhN5l3F+ofd0vj39zwQ== root@vi
  os7.vios.austin.ibx.com
$ mkauthkeys --add 'ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAvjDZ
  sS0guWzfzfp9BbweG0QMXv1tbDrtyWsgPbA2ExHA+xduWA51K0oFGarK2F
  C7e7NjKW+UmgQbrh/KSyKKwozjp4xWGNGhLmfan85ZpFR7wy9UQG1bLgXZ
  xYrY7yyQQQODjvwosWAfzkjpG3iW/xmWD5PKLBmob2QkKJbxjne+wqGwHT
  RYDGIiyhCBIdfFaLZgkXTZ2diZ98rL8LIv3qb+TsM1B28AL4t+1OGGeW24
  2lsB+8p4kamPJCYfKePHo67yP4NyKyPBFHY3TpTrca4/y1KEBT0Va3Pebr
  5JEIUvWYs6/RW+bUQk1Sb6eYbcRJFHhN5l3F+ofd0vj39zwQ== root@vi
  os7.vios.austin.ibx.com'
```

Results

When you are finished, you can view the data that is gathered by the monitoring agent from the Tivoli Enterprise Portal.

Related information

IBM Tivoli Monitoring version 6.2.1 documentation

Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide

Configuring the IBM Tivoli Usage and Accounting Manager agent

You can configure and start the IBM Tivoli Usage and Accounting Manager agent on the Virtual I/O Server.

About this task

With Virtual I/O Server 1.4, you can configure the IBM Tivoli Usage and Accounting Manager agent on the Virtual I/O Server. Tivoli Usage and Accounting Manager helps you track, allocate, and invoice your IT costs by collecting, analyzing, and reporting on the actual resources used by entities such as cost centers, departments, and users. Tivoli Usage and Accounting Manager can gather data from multi-tiered data centers that include Windows, AIX, Virtual I/O Server, HP/UX Sun Solaris, Linux, IBM i, and VMware.

Before you start, ensure that the Virtual I/O Server is installed. The Tivoli Usage and Accounting Manager agent is packaged with the Virtual I/O Server and is installed when the Virtual I/O Server is installed. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 90.

To configure and start the Tivoli Usage and Accounting Manager agent, complete the following steps:

Procedure

1. Optional: Add optional variables to the A_config.par file to enhance data collection.

   The A_config.par file is located at /home/padmin/tivoli/ituam/A_config.par. For more information about additional data collectors available for the ITUAM agent on the Virtual I/O Server, see the IBM Tivoli Usage and Accounting Manager Information Center.

2. List all of the available Tivoli Usage and Accounting Manager agents by using the **lssvc** command.

   For example,

   ```
   $lssvc
   ITUAM_base
   ```

3. Based on the output of the **lssvc** command, decide which Tivoli Usage and Accounting Manager agent you want to configure.
   For example, `ITUAM_base`
4. List all of the attributes that are associated with the Tivoli Usage and Accounting Manager agent by using the **cfgsvc** command.
   For example:

   ```
   $cfgsvc –ls ITUAM_base
    ACCT_DATA0
   ACCT_DATA1
   ISYSTEM
   IPROCESS
   ```

5. Configure the Tivoli Usage and Accounting Manager agent with its associated attributes by using the **cfgsvc** command:

   ```
   cfgsvc ITUAM_agent_name -attr ACCT_DATA0=value1 ACCT_DATA1=value2 ISYSTEM=value3
   IPROCESS=value4
   ```

   Where:

   - *ITUAM_agent_name* is the name of the Tivoli Usage and Accounting Manager agent. For example, ITUAM_base.
   - *value1* is the size (in MB) of the first data file that holds daily accounting information.
   - *value2* is the size (in MB) of the second data file that holds daily accounting information.
   - *value3* is the time (in minutes) when the agent generates system interval records.
   - *value4* is the time (in minutes) when the system generates aggregate process records.
6. Start the Tivoli Usage and Accounting Manager agent by using the **startsvc** command.
   For example:

   ```
   startsvc ITUAM_base
   ```

Results

After you start the Tivoli Usage and Accounting Manager agent, it begins to collect data and generate log files. You can configure the Tivoli Usage and Accounting Manager server to retrieve the log files, which are then processed by the Tivoli Usage and Accounting Manager Processing Engine. You can work with the data from the Tivoli Usage and Accounting Manager Processing Engine as follows:

- You can generate customized reports, spreadsheets, and graphs. Tivoli Usage and Accounting Manager provides full data access and reporting capabilities by integrating Microsoft SQL Server Reporting Services or Crystal Reports with a Database Management System (DBMS).
- You can view high-level and detailed cost and usage information.
- You can allocate, distribute, or charge IT costs to users, cost centers, and organizations in a manner that is fair, understandable, and reproducible.

For more information, see the IBM Tivoli Usage and Accounting Manager Information Center.

Related reference

Configuration attributes for IBM Tivoli agents and clients

Learn about required and optional configuration attributes and variables for the IBM Tivoli Monitoring agent, the IBM Tivoli Usage and Accounting Manager agent, the IBM Tivoli Storage Manager client, and the Tivoli Storage Productivity Center agents.

Configuring the IBM Tivoli Storage Manager client

You can configure theIBM Tivoli Storage Manager client on the Virtual I/O Server.

About this task

With Virtual I/O Server 1.4, you can configure the Tivoli Storage Manager client on the Virtual I/O Server. With Tivoli Storage Manager, you can protect your data from failures and other errors by storing backup and disaster-recovery data in a hierarchy of auxiliary storage. Tivoli Storage Manager can help protect computers that run various different operating environments, including the Virtual I/O Server, on various different hardware, including IBM Power Systems servers. If you configure the Tivoli Storage Manager client on the Virtual I/O Server, you can include the Virtual I/O Server in your standard backup framework.

Before you start, ensure that the Virtual I/O Server is installed. The Tivoli Storage Manager client is packaged with the Virtual I/O Server and is installed when the Virtual I/O Server is installed. For instructions, see “Installing the Virtual I/O Server and client logical partitions” on page 90.

To configure and start the Tivoli Storage Manager client, complete the following steps:

Procedure

1. List all of the available Tivoli Storage Manager clients by using the **lssvc** command.

   For example,

   ```
   $lssvc
   TSM_base
   ```

2. Based on the output of the **lssvc** command, decide which Tivoli Storage Manager client you want to configure.
   For example, `TSM_base`
3. List all of the attributes that are associated with the Tivoli Storage Manager client by using the **cfgsvc** command.
   For example:

   ```
   $cfgsvc -ls TSM_base
    SERVERNAME
   SERVERIP
   NODENAME
   ```

4. Configure the Tivoli Storage Manager client with its associated attributes by using the **cfgsvc** command:

   ```
   cfgsvc TSM_client_name -attr SERVERNAME=hostname SERVERIP=name_or_address NODENAME=vios
   ```

   Where:

   - *TSM_client_name* is the name of the Tivoli Storage Manager client. For example, TSM_base.
   - *hostname* is the host name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.
   - *name_or_address* is the IP address or domain name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.
   - *vios* is the name of the machine on which the Tivoli Storage Manager client is installed. The name must match the name that is registered on the Tivoli Storage Manager server.
5. Ask the Tivoli Storage Manager administrator to register the client node, the Virtual I/O Server, with the Tivoli Storage Manager server.

For more information about the IBM Tivoli Storage Manager, see the IBM Tivoli Storage Manager documentation.

Results

After you are finished, you are ready to back up and restore the Virtual I/O Server by using the Tivoli Storage Manager. For instructions, see the following procedures:

- "Backing up the Virtual I/O Server by using IBM Tivoli Storage Manager" on page 215
- "Restoring the Virtual I/O Server by using IBM Tivoli Storage Manager" on page 223

Configuring the Tivoli Storage Productivity Center agents

You can configure and start the Tivoli Storage Productivity Center agents on the Virtual I/O Server. Tivoli Storage Productivity Center is also known as IBM Tivoli Storage Productivity Center and IBM Spectrum Control.

About this task

With Virtual I/O Server 1.5.2, you can configure the Tivoli Storage Productivity Center agents on the Virtual I/O Server. Tivoli Storage Productivity Center is an integrated, storage infrastructure management suite that is designed to help simplify and automate the management of storage devices, storage networks, and capacity utilization of file systems and databases. When you configure the Tivoli Storage Productivity Center agents on the Virtual I/O Server, you can use the Tivoli Storage Productivity Center user interface to collect and view information about the Virtual I/O Server.

Note: The Tivoli Storage Productivity Center agent Version 6.2.2.0, or later, is included on the Virtual I/O Expansion media. This version of the Tivoli Storage Productivity Center agent requires the GSKit8 libraries, which are also included on the Virtual I/O Expansion media.

Before you start, complete the following tasks:

1. Use the **ioslevel** command to verify that the Virtual I/O Server is at Version 1.5.2, or later.
2. Ensure that there are no other operations running on the Virtual I/O Server. Configuring the Tivoli Storage Productivity Center consumes all of the processing time.
3. In addition to the memory required by the Virtual I/O Server logical partition, ensure that you have allocated a minimum of 1 GB of memory to the Virtual I/O Server for the Tivoli Storage Productivity Center agents.

To configure and start the Tivoli Storage Productivity Center agents, complete the following steps:

Procedure

1. List all of the available Tivoli Storage Productivity Center agents by using the **lssvc** command.

   For example,

   ```
   $lssvc
   TPC
   ```

   The Tivoli Storage Productivity Center agent includes both the TPC_data and TPC_fabric agents. When you configure the Tivoli Storage Productivity Center agent, you configure both the TPC_data and TPC_fabric agents.
2. List all of the attributes that are associated with the Tivoli Storage Productivity Center agent by using the **lssvc** command.
   For example:

   ```
   $lssvc TPC
   A:
   S:
   devAuth:
   caPass:
   caPort:
   ```

```
amRegPort:
amPubPort:
dataPort:
devPort:
newCA:
oldCA:
daScan:
daScript:
daInstall:
faInstall:
U:
```

The A, S, devAuth, and caPass attributes are required. The remainder of the attributes is optional. For more information about the attributes, see "Configuration attributes for IBM Tivoli agents and clients" on page 257.

3. Configure the Tivoli Storage Productivity Center agent with its associated attributes by using the **cfgsvc** command:

   ```
   cfgsvc TPC -attr S=tpc_server_hostname A=agent_manager_hostname devAuth=password_1
   caPass=password_2
   ```

   Where:

   - *tpc_server_hostname* is the host name or IP address of the Tivoli Storage Productivity Center server that is associated with the Tivoli Storage Productivity Center agent.
   - *agent_manager_hostname* is the name or IP address of the Agent Manager.
   - *password_1* is the password that is required to authenticate to the Tivoli Storage Productivity Center device server.
   - *password_2* is the password that is required to authenticate to the common agent.
4. Select the language that you want to use during the installation and configuration.
5. Accept the license agreement to install the agents according to the attributes specified in step "3" on page 191.
6. Start each Tivoli Storage Productivity Center agent by using the **startsvc** command:
   - To start the TPC_data agent, run the following command:

     ```
     startsvc TPC_data
     ```

   - To start the TPC_fabric agent, run the following command:

     ```
     startsvc TPC_fabric
     ```

Results

After you start the Tivoli Storage Productivity Center agents, you can perform the following tasks by using the Tivoli Storage Productivity Center user interface:

1. Run a discovery job for the agents on the Virtual I/O Server.
2. Run probes, scans, and ping jobs to collect storage information about the Virtual I/O Server.
3. Generate reports by using the Fabric Manager and the Data Manager to view the storage information gathered.
4. View the storage information gathered by using the topology Viewer.

For more information, see the *Tivoli Storage Productivity Center support for agents on a Virtual I/O Server* PDF file. To view or download the PDF file, go to the Planning for the Virtual I/O Server website.

Configuring the Virtual I/O Server as an LDAP client

Virtual I/O Server Version 1.4 can be configured as an LDAP client and then you can manage Virtual I/O Server from an LDAP server.

Before you begin

Before you start, gather the following information:

- The name of the Lightweight Directory Access Protocol (LDAP) server or servers to which you want the Virtual I/O Server to be an LDAP client.
- The administrator distinguish name (DN) and password for the LDAP server or servers to which you want the Virtual I/O Server to be an LDAP client.

About this task

To configure the Virtual I/O Server as an LDAP client, complete the following steps:

Procedure

1. Set up the LDAP client by running the following command:

   ```
   mkldap -host ldapserv1 -bind cn=admin -passwd adminpwd
   ```

   Where,

   - *ldapserv1* is the LDAP server or list of LDAP servers to which you want the Virtual I/O Server to be an LDAP client
   - *cn=admin* is the administrator DN of *ldapserv1*
   - *adminpwd* is the password for *cn=admin*

   Configuring the LDAP client automatically starts communication between the LDAP server and the LDAP client (the Virtual I/O Server). To stop communication, use the **stopnetsvc** command.
2. Change Virtual I/O Server users to LDAP users by running the following command:

   ```
   chuser -ldap -attr Attributes=Value username
   ```

   where, *username* is the name of the user you want to change to an LDAP user.

Configuring the Virtual I/O Server for the VSN capability

If you are using the Hardware Management Console (HMC) Version 7 Release 7.7.0, or later, you can use Virtual Station Interface (VSI) profiles with virtual Ethernet adapters in logical partitions and assign the Virtual Ethernet Port Aggregator (VEPA) switching mode to virtual Ethernet switches.

When you use the Virtual Ethernet Bridge (VEB) switching mode in virtual Ethernet switches, the traffic between logical partitions is not visible to the external switches. However, when you use the VEPA switching mode, the traffic between logical partitions is visible to the external switches. This visibility helps you to use features such as security that are supported by the advanced switching technology. Automated VSI discovery and configuration with the external Ethernet bridges simplifies the switch configuration for the virtual interfaces that are created with logical partitions. The profile-based VSI management policy definition provides flexibility during configuration and maximizes the benefits of automation.

The configuration requirements on the Virtual I/O Server (VIOS) to use the VSN capability follow:

- At least one VIOS logical partition that is servicing the virtual switch must be active and must support the VEPA switching mode.
- The external switches that are connected to the shared Ethernet adapter must support the VEPA switching mode.

- The **lldp** daemon must be running on the VIOS and must be managing the shared Ethernet adapter.
- From the VIOS command-line interface, run the **chdev** command to change the value of the *lldpsvc* attribute of the shared Ethernet adapter device to *yes*. The default value of the *lldpsvc* attribute is *no*. Run the **lldpsync** command to notify the change to the running **lldpd** daemon.

  Note: The *lldpsvc* attribute must be set to the default value before you remove the shared Ethernet adapter. Otherwise, removal of the shared Ethernet adapter fails.
- For redundancy shared Ethernet adapter setup, the trunk adapters might be attached to a virtual switch that is set to the VEPA mode. In this case, attach the control channel adapters of the shared Ethernet adapter to another virtual switch that is always set to the virtual Ethernet bridging (VEB) mode. The shared Ethernet adapter that is in the high availability mode does not work when the control channel adapter that is associated with the virtual switches is in the VEPA mode.

Restriction: To use VSN capability, you cannot configure a shared Ethernet adapter to use link aggregation or an Etherchannel device as the physical adapter.

Related information

Verifying that the server supports the virtual server network capability

Changing the virtual switch mode setting

Managing the Virtual I/O Server

You can manage virtual Small Computer Serial Interface (SCSI) and virtual Ethernet devices on the Virtual I/O Server, as well as backup, restore, update, and monitor the Virtual I/O Server.

Managing storage

You can import and export volume groups and storage pools, map virtual disks to physical disks, increase virtual Small Computer Serial Interface (SCSI) device capacity, change the virtual SCSI queue depth, back up and restore files and file systems, and collect and view information by using the Tivoli Storage Productivity Center.

Importing and exporting volume groups and logical volume storage pools

You can use the **importvg** and **exportvg** commands to move a user-defined volume group from one system to another.

About this task

Consider the following when you import and export volume groups and logical volume storage pools:

- The import procedure introduces the volume group to its new system.
- You can use the **importvg** command to reintroduce a volume group or logical volume storage pool to the system that it had been previously associated with and had been exported from.
- The **importvg** command changes the name of an imported logical volume if a logical volume of that name already exists on the new system. If the **importvg** command must rename a logical volume, it prints an error message to standard error.
- The export procedure removes the definition of a volume group from a system.
- You can use the **importvg** and **exportvg** commands to add a physical volume that contains data to a volume group by putting the disk to be added in its own volume group.
- The rootvg volume group cannot be exported or imported.

Importing volume groups and logical volume storage pools

You can use the **importvg** command to import a volume group or logical volume storage pool.

About this task

To import a volume group or logical volume storage pool, complete the following steps:

Procedure

1. Run the following command to import the volume group or logical volume storage pool:

   ```
   importvg -vg volumeGroupName physicalVolumeName
   ```

   Where:

   - *volumeGroupName* is an optional parameter that specifies the name to use for the imported volume group.
   - *physicalVolumeName* is the name of a physical volume that belongs to the imported volume group.
2. If you know that the imported volume group or logical volume storage pool is not the parent of the virtual media repository or any file storage pools, then you are finished importing the volume group or logical volume storage pool and do not need to complete the remaining steps.
3. If you know that imported volume group or logical volume storage pool is the parent of the virtual media repository or any file storage pools, or if you are unsure, then complete the following steps:

   a) Run the `mount all` command to mount any file systems that contain in the imported volume group or logical volume storage pool.

      This command might return errors for file systems that are already mounted.

   b) If you are importing a volume group or logical volume storage to the same system from which you exported it, run the `cfgdev` to reconfigure any devices that were unconfigured when you exported the volume group or logical volume storage pool.

What to do next

To export a volume group or logical volume storage pool, see "Exporting volume groups and logical volume storage pools" on page 194.

Exporting volume groups and logical volume storage pools

You can use the **exportvg** command to export a volume group or logical volume storage pool.

Before you begin

Before you start, complete the following tasks:

1. Determine whether the volume group or logical volume storage pool that you plan to export is a parent to the virtual media repository or to any file storage pools by completing the following steps:

   a. Run the **lsrep** command to determine whether the volume group or logical volume storage pool that you plan to export is a parent of the virtual media repository. The **Parent Pool** field displays the parent volume group or logical volume pool of the virtual media repository.

   b. Run the following command to determine whether a file storage pool is a child of the volume group or logical volume pool that you plan to export:

      ```
      lssp -detail -sp FilePoolName
      ```

      The results list the parent volume group or logical volume storage pool of the file storage pool.
2. If the volume group or logical volume storage pool that you plan to export is a parent of the virtual media repository or a file storage pool, then complete the following steps.

Table 42. Prerequisites steps if the volume group or logical volume storage pool is a parent of the virtual media repository or a file storage pool

Parent of Virtual Media Repository	**Parent of a file storage pool**

a. Unload the backing device of each file-backed optical virtual target device (VTD) that has a media file loaded, by completing the following steps:
i) Retrieve a list of the file-backed optical VTDs by running the following command:
`lsmap -all -type file_opt`
ii) For each device that shows a backing device, run the following command to unload the backing device:
`unloadopt -vtd` *`VirtualTargetDevice`*
b. Unmount the Virtual Media Repository file system by running the following command:
`unmount /var/vio/VMLibrary`
	a. Unconfigure the virtual target devices (VTDs) associated with the files that are contained in the file storage pools by completing the following steps:
	i) Retrieve a list of VTDs by running the following command:
	`lssp -bd -sp` *`FilePoolName`*
	where *FilePoolName* is the name of a file storage pool that is a child of the volume group or logical volume storage pool that you plan to export.
	ii) For each file that lists a VTD, run the following command:
	`rmdev -dev` *`VirtualTargetDevice`* `-ucfg`
	b. Unmount the file storage pool by running the following command:
	`unmount /var/vio/storagepools/`*`FilePoolName`*
	where *FilePoolName* is the name of the file storage pool to be unmounted.

About this task

To export the volume group or logical volume storage pool, run the following commands:

Procedure

1. `deactivatevg` *`VolumeGroupName`*
2. `exportvg` *`VolumeGroupName`*

   Where, *volumeGroupName* is an optional parameter that specifies the name to use for the imported volume group.

What to do next

To import a volume group or logical volume storage pool, see “Importing volume groups and logical volume storage pools” on page 193.

Mapping virtual disks to physical disks

Find instructions for mapping a virtual disk on a client logical partition to its physical disk on the Virtual I/O Server.

About this task

This procedure shows how to map a virtual Small Computer Serial Interface (SCSI) disk on an AIX or Linux client logical partition to the physical device (disk or logical volume) on the Virtual I/O Server.

To map a virtual disk to a physical disk, you need the following information. This information is gathered during this procedure:

- Virtual device name
- Slot number of the virtual SCSI client adapter
- Logical unit number (LUN) of the virtual SCSI device
- Client logical partition ID

Follow these steps to map a virtual disk on an AIX or Linux client logical partition to its physical disk on the Virtual I/O Server:

Procedure

1. Display virtual SCSI device information on the AIX or Linux client logical partition by typing the following command:

   ```
   lscfg -l devicename
   ```

   This command returns results similar to the following:

   ```
   U9117.570.1012A9F-V3-C2-T1-L810000000000  Virtual SCSI Disk Drive
   ```

2. Record the slot number, which is located in the output, following the card location label *C*. This identifies the slot number of the virtual SCSI client adapter. In this example, the slot number is 2.
3. Record the LUN, which is located in the output, following the LUN label *L*. In this example, the LUN is 810000000000.
4. Record the logical partition ID of the AIX or Linux client logical partition:

   a. Connect to the AIX or Linux client logical partition. For example, by using Telnet.

   b. On the AIX or Linux logical partition, run the `uname -L` command.

      Your results must look similar to the following:

      ```
      2  fumi02
      ```

      The logical partition ID is the first number listed. In this example, the logical partition ID is 2. This number is used in the next step.

   c. Type `exit`.

5. If you have multiple Virtual I/O Server logical partitions running on your system, determine which Virtual I/O Server logical partition is serving the virtual SCSI device. Use the slot number of the client adapter that is linked to a Virtual I/O Server, and a server adapter. Use the HMC command line to list information about virtual SCSI client adapters in the client logical partition.

   Log in to the HMC, and from the HMC command line, type `lshwres` . Specify the managed console name for the **-m** parameter and the client logical partition ID for the **lpar_ids** parameter.

   Note:

   - The managed console name, which is used for the **-m** parameter, is determined by typing `lssyscfg -r sys -F name` from the HMC command line.
   - Use the client logical partition ID recorded in Step 4 for the **-lpar_ids** parameter.

   For example:

   ```
   lshwres -r virtualio --rsubtype scsi -m fumi --filter lpar_ids=2
   ```

   This example returns results similar to the following:

   ```
   lpar_name=fumi02,lpar_id=2,slot_num=2,state=null,adapter_type=client,remote_lpar_id=1,
   remote_lpar_name=fumi01,remote_slot_num=2,is_required=1,backing_devices=none
   ```

Record the name of the Virtual I/O Server located in the **remote_lpar_name** field and slot number of the virtual SCSI server adapter, which is located in the **remote_slot_num=2** field. In this example, the name of the Virtual I/O Server is fumi01 and the slot number of the virtual SCSI server adapter is 2.

6. Log in to the Virtual I/O Server.
7. List virtual adapters and devices on the Virtual I/O Server by typing the following command:

   ```
   lsmap -all
   ```

8. Find the virtual SCSI server adapter (vhost*X*) that has a slot ID that matches the remote slot ID recorded in Step 5. On that adapter, run the following command:

   ```
   lsmap -vadapter devicename
   ```

9. From the list of devices, match the LUN recorded in step "3" on page 196 with LUNs listed. This is the physical device.

Increasing virtual SCSI device capacity

As storage demands increase for virtual client logical partitions, you can add physical storage to increase the size of your virtual devices and allocate that storage to your virtual environment.

About this task

You can increase the capacity of your virtual Small Computer Serial Interface (SCSI) devices by increasing the size of physical or logical volumes. With Virtual I/O Server Version 1.3 and later, you can do this without disrupting client operations. To increase the size of files and logical volumes based on storage pools, the Virtual I/O Server must be at Version 1.5 or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.

Tip: If you are using the HMC, Version 7 release 3.4.2, or later, you can use the HMC graphical interface to increase the capacity of a virtual SCSI device on aVirtual I/O Server.

To increase virtual SCSI device capacity, complete the following steps:

Procedure

1. Increase the size of the physical volumes, logical volumes, or files:

   - Physical volumes: Consult your storage documentation to determine whether your storage subsystem supports expanding the size of a logical unit number (LUN). For Virtual I/O Server Version 2.1.2.0, ensure that the Virtual I/O Server recognizes and adjusts to the new size by running the following command: `chvg -chksize` *vg1*, where *vg1* is the name of the expanding volume group.

     The Virtual I/O Server examines all the disks in volume group *vg1* to determine whether they have grown in size. For those disks that have grown in size, the Virtual I/O Server attempts to add additional physical partitions to the physical volumes. If necessary, the Virtual I/O Server determines the correct 1016 multiplier and conversion to a large volume group.

   - Logical volumes based on volume groups:

     a. Run the **extendlv** command. For example: `extendlv lv3 100M`. This example increases logical volume *lv3* by 100 MB.

     b. If there is no additional space in the logical volume, complete the following tasks:

        i) Increase the size of the volume group by completing one of the following steps:

           - Increase the size of the physical volumes. Consult your storage documentation for instructions.
           - Add physical volumes to a volume group by running the **extendvg** command. For example: `extendvg vg1 hdisk2`. This example adds physical volume *hdisk2* to volume group *vg1*.

        ii) Allocate the increased volume to partitions by resizing logical volumes. Run the **extendlv** command to increase the size of a logical volume.

- Logical volumes based on storage pools:

    a. Run the **chbdsp** command. For example:`chbdsp -sp lvPool -bd lv3 -size 100M`. This example increases logical volume *lv3* by 100 MB.

    b. If there is no additional space in the logical volume, complete the following tasks:

        i) Increase the size of the logical volume storage pool by completing one of the following steps:

        - Increase the size of the physical volumes. Consult your storage documentation for instructions.
        - Add physical volumes to the storage pool by running the **chsp** command. For example: `chsp -add -sp sp1 hdisk2`. This example adds physical volume *hdisk2* to storage pool *sp1*.

        ii) Allocate the increased volume to partitions by resizing logical volumes. Run the **chbdsp** command to increase the size of a logical volume.

- Files:

    a. Run the **chbdsp** command. For example:`chbdsp -sp fbPool -bd fb3 -size 100M`. This example increases file *fb3* by 100 MB.

    b. If there is no additional space in the file, increase the size of the file storage pool by running the **chsp** command. For example:`chsp -add -sp fbPool -size 100M`. This example increases file storage pool *fbPool* by 100 MB.

    c. If there is no additional space in the file storage pool, increase the size of the parent storage pool by completing one of the following tasks:

    - Increase the size of the physical volumes. Consult your storage documentation for instructions.
    - Add physical volumes to the parent storage pool by running the **chsp** command. For example:`chsp -add -sp sp1 hdisk2`. This example adds physical volume *hdisk2* to storage pool *sp1*.
    - Increase the size of the file storage pool by running the **chsp** command.

2. If you are running Virtual I/O Server versions before 1.3, then you need to either reconfigure the virtual device (by using the **cfgdev** command) or restart the Virtual I/O Server.
3. If you are running Virtual I/O Server Version 1.3 or later, then restarting or reconfiguring a logical partition is not required to begin by using the additional resources. If the physical storage resources have been set up and properly allocated to the system as a system resource, as soon as the Virtual I/O Server recognizes the changes in storage volume, the increased storage capacity is available to the client logical partitions.
4. On the client logical partition, ensure that the operating system recognizes and adjusts to the new size. For example, if AIX is the operating system on the client logical partition, run the following command:

    ```
    chvg -g vg1
    ```

    In this example, AIX examines all the disks in volume group *vg1* to see if they have grown in size. For the disks that have grown in size, AIX attempts to add additional physical partitions to physical volumes. If necessary, AIX determines proper 1016 multiplier and conversion to the large volume group.
    For example, if AIX is the operating system on the client logical partition, run the following command:

    ```
    chvg -g vg1
    ```

    In this example, AIX examines all the disks in volume group *vg1* to see if they have grown in size. For the disks that have grown in size, AIX attempts to add additional physical partitions to physical volumes. If necessary, AIX determines proper 1016 multiplier and conversion to the large volume group.

Related information

chlv Command

chvg Command

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Changing a storage pool for a VIOS logical partition using the HMC

Changing the virtual SCSI queue depth

Increasing the virtual Small Computer Serial Interface (SCSI) queue depth might provide performance improvements for some virtual configurations. Understand the factors that are involved in determining a change to the virtual SCSI queue depth value.

The virtual SCSI queue depth value determines how many requests the disk head driver queues to the virtual SCSI client driver at any one time. For AIX client logical partitions, you can change this value from the default value of 3 to a value in the range 1 - 256 by using the **chdev** command. For Linux client logical partitions, you can change this value from the default value of 16 to a value in the range 1 - 256 by using the **echo** command. For IBM i client logical partitions, the queue depth value is 32 and cannot be changed.

Increasing this value might improve the throughput of the disk in specific configurations. However, several factors must be considered. These factors include the value of the queue-depth attribute for all of the physical storage devices on the Virtual I/O Server being used as a virtual target device by the disk instance on the client logical partition, and the maximum transfer size for the virtual SCSI client adapter instance that is the parent device for the disk instance.

For AIX and Linux client logical partitions, the maximum transfer size for virtual SCSI client adapters is set by the Virtual I/O Server, which determines the value based on the resources available on the server and the maximum transfer size set for the physical storage devices on that server. Other factors include the queue depth and maximum transfer size of other devices that are involved in mirrored-volume-group or Multipath I/O (MPIO) configurations. Increasing the queue depth for some devices might reduce the resources available for other devices on that same shared adapter and decrease the throughput for those devices. For IBM i client logical partitions, the queue depth value is 32 and cannot be changed.

To change the queue depth for an AIX client logical partition, on the client logical partition, use the **chdev** command with the **queue_depth=value** attribute as in the following example:

```
chdev -l hdiskN -a "queue_depth=value"
```

hdiskN represents the name of a physical volume and *value* is the value that you assign in the range 1 - 256.

To change the queue depth for a Linux client logical partition on the client logical partition, use the **echo** command as in the following example:

```
echo 16 > /sys/devices/vio/30000003/host0/target0:0:1/0:0:1:0/queue_depth
```

By default, the value of the **queue_depth** attribute for a disk on the Linux operating system is 16.

To view the current setting for the queue_depth value, from the client logical partition issue the following command:

```
lsattr -El hdiskN
```

Backing up and restoring files and file systems

You can use the **backup** and **restore** commands to back up and restore individual files or entire file systems.

About this task

Backing up and restoring files and files systems can be useful for tasks, such as saving IBM i to physical tape or saving a file-backed device.

The following commands are used to back up and restore files and files systems.

Table 43. Backup and restore commands and their descriptions

Command	Description
backup	Backs up files and file systems to media, such as physical tape and disk. For example: • You can back up all the files and subdirectories in a directory by using full path names or relative path names. • You can back up the root file system. • You can back up all the files in the root file system that have been modified since the last backup. • You can back up virtual optical media files from the virtual media repository.
restore	Reads archives created by the **backup** command and extracts the files that are stored there. For example: • You can restore a specific file into the current directory. • You can restore a specific file from tape into the virtual media repository. • You can restore a specific directory and the contents of that directory from a file name archive or a file system archive. • You can restore an entire file system. • You can restore only the permissions or only the ACL attributes of the files from the archive.

Managing storage by using the Tivoli Storage Productivity Center

You can use the Tivoli Storage Productivity Center collect and view information about the Virtual I/O Server.

About this task

With Virtual I/O Server 1.5.2, you can install and configure the Tivoli Storage Productivity Center agents on the Virtual I/O Server. Tivoli Storage Productivity Center is an integrated, infrastructure management suite for storage that is designed to help simplify and automate the management of storage devices, storage networks, and capacity utilization of file systems and databases. When you install and configure the Tivoli Storage Productivity Center agents on the Virtual I/O Server, you can use the Tivoli Storage Productivity Center interface to collect and view information about the Virtual I/O Server. You can then perform the following tasks by using the Tivoli Storage Productivity Center interface:

1. Run a discovery job for the agents on the Virtual I/O Server.
2. Run probes, run scans, and ping jobs to collect storage information about the Virtual I/O Server.
3. Generate reports by using the Fabric Manager and the Data Manager to view the storage information gathered.
4. View the storage information gathered by using the topology Viewer.

Related tasks

Configuring the Tivoli Storage Productivity Center agents
You can configure and start the Tivoli Storage Productivity Center agents on the Virtual I/O Server. Tivoli Storage Productivity Center is also known as IBM Tivoli Storage Productivity Center and IBM Spectrum Control.

Managing networks

You can change the network configuration of the Virtual I/O Server logical partition, enable and disable GARP VLAN Registration Protocol (GVRP) on your **Shared Ethernet Adapters**, use Simple Network Management Protocol (SNMP) to manage systems and devices in complex networks, and upgrade to Internet Protocol version 6 (IPv6).

Removing the network configuration of the Virtual I/O Server logical partition

You can remove the network settings on the Virtual I/O Server (VIOS) logical partition.

About this task

The following list describes how to remove the network settings on the VIOS partition:

Procedure

- To remove the configuration from a network interface, type the following command:

  ```
  rmtcpip [-interface interface]
  ```

- To remove only Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) from an interface, type the following command:

  ```
  rmtcpip [-interface interface] [-family family]
  ```

- To remove the IP configuration from the system, type the following command:

  ```
  rmtcpip -all
  ```

Results

Note: You cannot remove the IP configuration that is used for communication in a shared storage pool.

Dynamically adding or removing VLANs on the Virtual I/O Server

With the Virtual I/O Server Version 2.2, or later, you can add, change, or remove the existing set of VLANs for a virtual Ethernet adapter that is assigned to an active partition on a POWER7, POWER8, or POWER9 processor-based servers by using the Hardware Management Console (HMC).

Before you begin

Before you perform this task, ensure that you meet the following requirements:

- The server must be a POWER7, POWER8, or POWER9 processor-based servers, or later.
- The server firmware level must be at least AH720_064+ for high end servers, AM720_064+ for midrange servers, and AL720_064+ for low end servers.

  Note: The AL720_064+ server firmware level is only supported on POWER7 processor-based servers, or later.
- The Virtual I/O Server must be at Version 2.2, or later.
- The HMC must be at Version 7.7.2.0, with mandatory fix MH01235, or later.

About this task

You can use the HMC graphical interface or the **chhwres** command from the HMC command-line interface to add, remove, or modify VLANs for a virtual Ethernet adapter that is assigned to an active partition. You can also edit the IEEE standard of the virtual Ethernet adapter dynamically. To specify additional VLANs, you must set the virtual Ethernet adapter to the IEEE 802.1Q standard.

To add, remove, or modify VLANs on the Virtual I/O Server, complete the following steps:

Procedure

1. Run the **lssyscfg** command to verify if the managed system supports adding, removing, or modifying VLANs on the Virtual I/O Server. For example,

   ```
   lssyscfg -r sys -m <managed system> -F capabilities
   ```

   If the managed server supports adding, removing, or modifying VLANs, this command returns the `virtual_eth_dlpar_capable` value.
2. Use the **chhwres** command to add, remove, or modify additional VLANs to the virtual Ethernet adapter that is assigned to an active partition. You can also edit the IEEE standard of the virtual Ethernet adapter dynamically by using the **chhwres** command. For example,

   In this example, the VLAN ID 5 is added to the existing VLAN IDs for the virtual Ethernet adapter, and the virtual Ethernet adapter is set to the IEEE 802.1Q standard.

   ```
   chhwres -r virtualio --rsubtype eth -m <managed system> -o s     {-p <partition name> |
   --id <partition ID>} -s <virtual slot number> -a "addl_vlan_ids+=5,ieee_virtual_eth=1"
   ```

   In this example, the VLAN ID 6 is removed from the existing VLAN IDs for the virtual Ethernet adapter.

   ```
   chhwres -r virtualio --rsubtype eth -m <managed system> -o s     {-p <partition name> |
   --id <partition ID>} -s <virtual slot number> -a "addl_vlan_ids-=6"
   ```

   In this example, the VLAN IDs 2, 3, and 5 are assigned to the virtual Ethernet adapter instead of the existing VLAN IDs.

   ```
   chhwres -r virtualio --rsubtype eth -m <managed system> -o s     {-p <partition name> |
   --id <partition ID>} -s <virtual slot number> -a "addl_vlan_ids=2,3,5"
   ```

   You can provide a comma-separated list of VLANs to the attributes, **addl_vlan_ids=**, **addl_vlan_ids+=**, and **addl_vlan_ids-=**.
3. Use the **lshwres** command to query the virtual Ethernet adapter.

   ```
   lshwres -m <server> -r virtualio --rsubtype eth --level lpar -F
   ```

Enabling or disabling the virtual Ethernet adapter

You can remove the selected partition from the network by disabling the virtual Ethernet adapter (VEA) that is configured on the partition and then connecting it back to network by enabling that virtual Ethernet adapter.

Before you begin

Note: You must check whether enabling, disabling, or querying the VEA is supported.

By default, the virtual Ethernet adapter is enabled.

Procedure

1. To check whether enabling, disabling, or querying the VEA is supported, type the following command:

```
lssyscfg -r sys -F capabilities
```

The system displays the output as follows:

`virtual_eth_disable_capable`

Note: If the output is displayed as **`virtual_eth_disable_capable`**, enabling, disabling, or querying the VEA is supported.

2. To query the VEA, type the following command:

```
lshwres -m <server> -r virtualio --rsubtype eth --level lpar -F
```

3. To enable or disable the VEA, type the following command:

```
chhwres -m <server> -r virtualio --rsubtype eth -o {d | e} {-p <lpar name>
--id <lpar ID>} -s <slot number>
```

The description of flags is as follows:

- *d* - Disables the VEA.
- *e* - Enables the VEA

Note: The VEA can be disabled only when the logical partition capabilities support disabling of VEA. To disable VEA, the logical partition can either be in *Activated*, *Open Firmware*, or *Not Activated* state.

Enabling and disabling GVRP

You can enable and disable GARP VLAN Registration Protocol (GVRP) on your **Shared Ethernet Adapters** to control dynamic registration of VLANs over networks.

Before you begin

With Virtual I/O Server Version 1.4, **Shared Ethernet Adapters** support GARP VLAN Registration Protocol (GVRP) which is based on GARP (Generic Attribute Registration Protocol). GVRP allows for the dynamic registration of VLANs over networks.

By default, GVRP is disabled on **Shared Ethernet Adapters**.

Before you start, create and configure the Shared Ethernet Adapter. For instructions, see "Creating a virtual Ethernet adapter with the HMC Version 7 graphical interface" on page 178.

Procedure

To enable or disable GVRP, run the following command:

```
chdev -dev Name -attr gvrp=yes/no
```

Where:

- *Name* is the name of the Shared Ethernet Adapter.
- *yes/no* defines whether GVRP is enabled or disabled. Type `yes` to enable GVRP and type no to disable GVRP.

Managing SNMP on the Virtual I/O Server

Find commands for enabling, disabling, and working with SNMP on the Virtual I/O Server.

About this task

Simple Network Management Protocol (SNMP) is a set of protocols for monitoring systems and devices in complex networks. SNMP network management is based on the familiar client/server model that is widely used in Internet Protocol (IP) network applications. Each managed host runs a process called an agent.

The agent is a server process that maintains information about managed devices in the Management Information Base (MIB) database for the host. Hosts that are involved in network management decision-making can run a process called a manager. A manager is a client application that generates requests for MIB information and processes responses. In addition, a manager might send requests to agent servers to modify MIB information.

In general, network administrators use SNMP to more easily manage their networks for the following reasons:

- It hides the underlying system network.
- The administrator can manage and monitor all network components from one console.

SNMP is available on Virtual I/O Server Version 1.4 and later.

The following table lists the SNMP management tasks available on the Virtual I/O Server, as well as the commands you need to run to accomplish each task.

Table 44. Commands for working with SNMP on the Virtual I/O Server

Command	Task
startnetsvc	Enable SNMP.
snmpv3_ssw	Select which SNMP agent you want to run.
cl_snmp	Issue SNMP requests to agents.
cl_snmp	Process SNMP responses that are returned by agents.
snmp_info	Request MIB information that is managed by an SNMP agent.
snmp_info	Modify MIB information that is managed by an SNMP agent.
snmp_trap	Generate a notification, or trap, that reports an event to the SNMP manager with a specified message.
stopnetsvc	Disable SNMP.

Related information

Network Management

Configuring IPv6 on the Virtual I/O Server

To take advantage of enhancements, such as expanded addressing and routing simplification, use the **mktcpip** command to configure Internet Protocol version 6 (IPv6) on the Virtual I/O Server (VIOS).

About this task

IPv6 is the next generation of Internet Protocol and is gradually replacing the current Internet standard, Internet Protocol version 4 (IPv4). The key IPv6 enhancement is the expansion of the IP address space 32 - 128 bits, providing virtually unlimited, unique IP addresses. IPv6 provides several advantages over IPv4 including expanded routing and addressing, routing simplification, header format simplification, improved traffic control, autoconfiguration, and security.

Procedure

To configure IPv6 on the VIOS, type the following command:

```
mktcpip –auto [-interface interface] [-hostname hostname]
```

Where:

- `interface` specifies which interface you want to configure for IPv6.
- `hostname` specifies the host name of the system to be set.

This command automatically performs the following tasks:

- Configures IPv6 link-local addresses on the interfaces that are currently configured with IPv4.
- Starts the ndpd-host daemon.
- Ensures that the IPv6 configuration remains intact after you reboot the VIOS.

Note: You can also use the following command to configure static IPv6 address on a VIOS. However, IPv6 stateless autoconfiguration is suggested.

```
mktcpip -hostname HostName -inetaddr Address -interface Interface
[-start] [-plen PrefixLength] [-cabletype CableType] [-gateway Gateway]
[-nsrvaddr NameServerAddress -nsrvdomain Domain]
```

What to do next

If you decide that you want to undo the IPv6 configuration, run the **`rmtcpip`** command with the `-family` option. For instructions, see "Removing the network configuration of the Virtual I/O Server logical partition" on page 201.

Subscribing to Virtual I/O Server product updates

A subscription service is available for Virtual I/O Server users to stay current on news and product updates.

About this task

To subscribe to this service, follow these steps:

Procedure

1. Go to the IBM Support website.
2. In the **My Notifications** page, enter the product details in the **Product lookup** field and click **Subscribe**.
3. In the **Select document types** window, select the types of documents for which you want to receive notifications.
4. Click **Submit** to save the changes. Alternatively, you can click **Close** to cancel the changes and to close the window.

What to do next

After subscribing, you will be notified of all Virtual I/O Server news and product updates.

Updating the Virtual I/O Server

To install an update to the Virtual I/O Server, you can obtain the update either from a CD that contains the update or download the update.

About this task

To update the Virtual I/O Server, follow these steps:

Procedure

1. Make a backup of the Virtual I/O Server by following the steps in Backing up the Virtual I/O Server.

2. Download the required updates from the Fix Central website. Alternatively, you can get the updates from the update CD.
3. Install the update by using the **updateios** command. For example, if your update file set is located in the /home/padmin/update directory, type the following:

   ```
   updateios -install -accept -dev /home/padmin/update
   ```

   Notes:

   - The **updateios** command installs all updates that are located in the specified directory.
   - The Virtual I/O Server (VIOS) Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, does not support the -reject option of the **updateios** command.
   - To perform Live Partition Mobility after you install an update to the VIOS, ensure that you restart the Hardware Management Console (HMC).

Backing up the Virtual I/O Server

You can back up the Virtual I/O Server (VIOS) and user-defined virtual devices by using the **backupios** command or the **viosbr** command. You can also use IBM Tivoli Storage Manager to schedule backups and to store backups on another server.

About this task

The VIOS contains the following types of information that you need to back up: the VIOS itself and user-defined virtual devices.

- The VIOS includes the base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata. All this information is backed up when you use the **backupios** command.
- User-defined virtual devices include metadata, like virtual devices mappings, that define the relationship between the physical environment and the virtual environment. You can back up user-defined virtual devices in one of the following ways:
  - You can back up user-defined virtual devices by using the **viosbr** command. Use this option in situations where you plan to restore the configuration information to the same VIOS partition from which it was backed up.
  - You can back up user-defined virtual devices by saving the data to a location that is automatically backed up when you use the **backupios** command to back up the VIOS. Use this option in situations where you plan to restore the VIOS to a new or different system. (For example, in the event of a system failure or disaster.) Furthermore, in these situations, you must also back up the following components of your environment. Back up these components to fully recover your VIOS configuration:
    - External device configurations, such as storage area network (SAN) devices.
    - Resources that are defined on the Hardware Management Console (HMC), such as processor and memory allocations. In other words, back up your HMC partition profile data for the VIOS and its client partitions.
    - The operating systems and applications that run in the client logical partitions.

You can back up and restore the VIOS as follows.

Table 45. Backup and restoration methods for the VIOS

Backup method	Media	Restoration method
To tape	Tape	From tape
To DVD	DVD-RAM	From DVD

Table 45. Backup and restoration methods for the VIOS (continued)

Backup method	Media	Restoration method
To remote file system	`nim_resources.tar` image	From an HMC by using the Network Installation Management (NIM) on Linux facility and the **`installios`** command
To remote file system	`mksysb` image	From an AIX 5L NIM server and a standard `mksysb` system installation
To remote file system	`mksysb` image	From an AIX 5L NIM server and a standard `mksysb` system installation
Tivoli Storage Manager	`mksysb` image	Tivoli Storage Manager

Related tasks

Restoring the Virtual I/O Server
You can restore the Virtual I/O Server (VIOS) and user-defined virtual devices by using the **`installios`** command, the **`viosbr`** command, or IBM Tivoli Storage Manager.

Related information

backupios command

viosbr command

Backing up the Virtual I/O Server to tape

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to tape.

About this task

To back up the Virtual I/O Server to tape, follow these steps:

Procedure

1. Assign a tape drive to the Virtual I/O Server.
2. Get the device name by typing the following command:

   ```
   lsdev -type tape
   ```

   If the tape device is in the `Defined` state, type the following command, where *dev* is the name of your tape device:

   ```
   cfgdev -dev dev
   ```

3. Type the following command, where *tape_device* is the name of the tape device you want to back up to:

   ```
   backupios -tape tape_device
   ```

   This command creates a bootable tape that you can use to restore the Virtual I/O Server.
4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, you need to back up the user-defined virtual devices.

   For instructions, see "Backing up user-defined virtual devices by using the backupios command" on page 210.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Backing up the Virtual I/O Server to one or more DVDs

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to DVD.

About this task

To back up the Virtual I/O Server to one or more DVDs, follow these steps. Only DVD-RAM media can be used to back up the Virtual I/O Server.

Note: Vendor disk drives might support burning to additional disk types, such as CD-RW and DVD-R. Refer to the documentation for your drive to determine which disk types are supported.

Procedure

1. Assign an optical drive to the Virtual I/O Server logical partition.
2. Get the device name by typing the following command:

   ```
   lsdev -type optical
   ```

   If the device is in the `Defined` state, type:

   ```
   cfgdev -dev dev
   ```

3. Run the **backupios** command with the **-cd** option. Specify the path to the device. For example:

   ```
   backupios -cd /dev/cd0
   ```

   Note: If the Virtual I/O Server does not fit on one DVD, then the **backupios** command provides instructions for disk replacement and removal until all the volumes have been created.

   This command creates one or more bootable DVDs that you can use to restore the Virtual I/O Server.
4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices.

   For instructions, see "Backing up user-defined virtual devices by using the backupios command" on page 210.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Backing up the Virtual I/O Server to a remote file system by creating a `nim_resources.tar` file

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to a remote file system by creating a `nim_resources.tar` file.

Before you begin

Backing up the Virtual I/O Server to a remote file system creates the `nim_resources.tar` image in the directory you specify. The `nim_resources.tar` file contains all the necessary resources to restore the Virtual I/O Server, including the mksysb image, the bosinst.data file, the network boot image, and Shared Product Object Tree (SPOT) resource.

The **backupios** command empties the target_disks_stanza section of bosinst.data and sets `RECOVER_DEVICES=Default`. This allows the mksysb file that is generated by the command to be cloned to another logical partition. If you plan to use the `nim_resources.tar` image to install to a specific disk, then you need to repopulate the target_disk_stanza section of bosinst.data and replace this file in the `nim_resources.tar` image. All other parts of the nim_resources.tar image must remain unchanged.

Before you start, complete the following tasks:

1. Ensure that the remote file system is available and mounted.
2. Ensure that the Virtual I/O Server has root write access to the server on which the backup is to be created.

About this task

To back up the Virtual I/O Server to a remote file system, follow these steps:

Procedure

1. Create a mount directory where the backup image, `nim_resources.tar`, is to be written. For example, to create the directory `/home/backup`, type:

   ```
   mkdir /home/backup
   ```

2. Mount an exported directory on the mount directory. For example:

   ```
   mount server1:/export/ios_backup /home/backup
   ```

3. Run the **backupios** command with the **-file** option. Specify the path to the mounted directory. For example:

   ```
   backupios -file /home/backup
   ```

   This command creates a `nim_resources.tar` file that you can use to restore the Virtual I/O Server from the HMC.

4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices.

   For instructions, see "Backing up user-defined virtual devices by using the backupios command" on page 210.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Backing up the Virtual I/O Server to a remote file system by creating an mksysb image

You can back up the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata to a remote file system by creating an mksysb file.

Before you begin

Backing up the Virtual I/O Server to a remote file system creates the mksysb image in the directory you specify. The mksysb image is an installable image of the root volume group in a file.

Before you start, complete the following tasks:

1. If you plan to restore the Virtual I/O Server from a Network Installation Management (NIM) server, verify that the NIM server is at the latest release of AIX. To find the latest updates, see the Fix Central website.
2. If you plan to restore the Virtual I/O Server from a Network Installation Management (NIM) server, verify that the NIM server is at the latest release of AIX. To find the latest updates, see the Fix Central website.
3. Ensure that the remote file system is available and mounted.

4. Ensure that the Virtual I/O Server has root write access to the server on which the backup is created.

About this task

To back up the Virtual I/O Server to a remote file system, follow these steps:

Procedure

1. Create a mount directory where the backup image, mksysb image, is to be written. For example, to create the directory `/home/backup`, type:

   ```
   mkdir /home/backup
   ```

2. Mount an exported directory on the mount directory. For example:

   ```
   mount server1:/export/ios_backup /home/backup
   ```

   where, *server1* is the NIM server from which you plan to restore the Virtual I/O Server.
3. Run the **backupios** command with the **-file** option. Specify the path to the mounted directory. For example:

   ```
   backupios -file /home/backup/filename.mksysb -mksysb
   ```

   where, *filename* is the name of mksysb image that this command creates in the specified directory.

   You can use the mksysb image to restore the Virtual I/O Server from a NIM server.
4. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then you need to back up the user-defined virtual devices.

   For instructions, see "Backing up user-defined virtual devices by using the backupios command" on page 210.

Backing up user-defined virtual devices

You can back up user-defined virtual devices by saving the data to a location that is automatically backed up when you use the **backupios** command to back up the Virtual I/O Server (VIOS). Alternatively, you can back up user-defined virtual devices by using the **viosbr** command.

About this task

User-defined virtual devices include metadata, such as virtual device mappings, that define the relationship between the physical environment and the virtual environment. You can back up user-defined virtual devices in one of the following ways:

- You can back up user-defined virtual devices by saving the data to a location that is automatically backed up when you use the **backupios** command to back up the VIOS. Use this option in situations where you plan to restore the VIOS to a new or different system. (For example, in the event of a system failure or disaster.)
- You can back up user-defined virtual devices by using the **viosbr** command. Use this option in situations where you plan to restore the configuration information to the same VIOS partition from which it was backed up.

Related tasks

Restoring user-defined virtual devices
You can restore user-defined virtual devices on the Virtual I/O Server (VIOS) by restoring volume groups and manually re-creating virtual device mappings. Alternatively, you can restore user-defined virtual devices by using the **viosbr** command.

Backing up user-defined virtual devices by using the backupios command

In addition to backing up the Virtual I/O Server (VIOS), you must back up user-defined virtual devices (such as virtual device mappings) in case you have a system failure or disaster. In this situation, back up user-defined virtual devices by saving the data to a location that is automatically backed up when you use the **backupios** command to back up the VIOS.

Before you begin

User-defined virtual devices include metadata, such as virtual device mappings, that define the relationship between the physical environment and the virtual environment. In situations where you plan to restore the VIOS to a new or different system, you need to back up both the VIOS and user-defined virtual devices. (For example, in the event of a system failure or disaster.)

Before you start, complete the following tasks:

1. Back up the VIOS to tape, DVD, or a remote file system. For instructions, see one of the following procedures:
   - "Backing up the Virtual I/O Server to tape" on page 207
   - "Backing up the Virtual I/O Server to one or more DVDs" on page 208
   - "Backing up the Virtual I/O Server to a remote file system by creating a nim_resources.tar file" on page 208
   - "Backing up the Virtual I/O Server to a remote file system by creating an mksysb image" on page 209
2. Decide whether you want to create a script of the following procedure. Scripting these commands makes it easy to schedule automated backups of the information.

About this task

To back up user-defined virtual devices, complete the following steps:

Procedure

1. List volume groups (and storage pools) to determine what user-defined disk structures you want to back up by running the following command:

   ```
   lsvg
   ```

2. Activate each volume group (and storage pool) that you want to back up by running the following command for each volume group:

   ```
   activatevg volume_group
   ```

   where, *volume_group* is the name of the volume group (or storage pool) that you want to activate.
3. Back up each volume group (and storage pool) by running the following command for each volume group:

   ```
   savevgstruct volume_group
   ```

   where, *volume_group* is the name of the volume group (or storage pool) that you want to back up.

   This command writes a backup of the structure of a volume group (and therefore, a storage pool) to the **/home/ios/vgbackups** directory.
4. Save the information about network settings, adapters, users, and security settings to the `/home/padmin` directory by running each command with the **tee** command as follows:

   ```
   command | tee /home/padmin/filename
   ```

   Where,

   - *command* is the command that produces the information you want to save.
   - *filename* is the name of the file to which you want to save the information.

Table 46. Commands that provide the information to save

Command	**Description**
`cfgnamesrv -ls`	Shows all system configuration database entries that are related to domain name server information used by local resolver routines.
`entstat -all` *`devicename`* *devicename* is the name of a device whose attributes or statistics you want to save. Run this command for each device whose attributes or statistics you want to save.	Shows Ethernet driver and device statistics for the device specified.
`hostmap -ls`	Shows all entries in the system configuration database.
`ioslevel`	Shows the current maintenance level of the Virtual I/O Server.
`lsdev -dev` *`devicename`* `-attr` *devicename* is the name of a device whose attributes or statistics you want to save. Run this command for each device whose attributes or statistics you want to save.	Shows the attributes of the device specified.
`lsdev -type adapter`	Shows information about physical and logical adapters.
`lsuser`	Shows a list of all attributes of all the system users.
`netstat -routinfo`	Shows the routing tables, including the user-configured and current costs of each route.
`netstat -state`	Shows the state of all configured interfaces.
`optimizenet -list`	Shows characteristics of all network tuning parameters, including the current and reboot value, range, unit, type, and dependencies.
`viosecure -firewall view`	Shows a list of allowed ports.
`viosecure -view -nonint`	Shows all the security level settings for noninteractive mode.

Related tasks

Scheduling backups of the Virtual I/O Server and user-defined virtual devices by creating a script and crontab file entry
You can schedule regular backups of the Virtual I/O Server (VIOS) and user-defined virtual devices to ensure that your backup copy accurately reflects the current configuration.

Backing up user-defined virtual devices by using the viosbr command
You can back up user-defined virtual devices by using the **`viosbr`** command. Use the **`viosbr`** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Backing up user-defined virtual devices by using the `viosbr` command

You can back up user-defined virtual devices by using the **viosbr** command. Use the **viosbr** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Before you begin

You can use the **viosbr** command to back up all the relevant data to recover a VIOS after an installation. The **viosbr** command backs up all the device properties and the virtual devices configuration on the VIOS. You can include information about some or all of the following devices in the backup:

- Logical devices, such as storage pools, clusters, file-backed storage pools, the virtual media repository, and paging space devices.
- Virtual devices, such as Etherchannel, Shared Ethernet Adapter, virtual server adapters, and virtual-server Fibre Channel adapters.
- Device attributes for devices like disks, optical devices, tape devices, fscsi controllers, Ethernet adapters, Ethernet interfaces, and logical **Host Ethernet Adapters**.

Before you start, run the **ioslevel** command to verify that the VIOS is at Version 2.1.2.0, or later.

Procedure

To back up all the device attributes and logical and virtual device mappings on the VIOS, run the following command:

```
viosbr -backup –file /tmp/myserverbackup
```

where, */tmp/myserverbackup* is the file to which you want to back up the configuration information.

Related tasks

Restoring user-defined virtual devices by using the viosbr command
You can restore user-defined virtual devices by using the **viosbr** command. Use the **viosbr** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Scheduling backups of user-defined virtual devices by using the viosbr command
You can schedule regular backups of the user-defined virtual devices on the Virtual I/O Server (VIOS) logical partition. Scheduling regular backups ensures that your backup copy accurately reflects the current configuration.

Backing up user-defined virtual devices by using the backupios command
In addition to backing up the Virtual I/O Server (VIOS), you must back up user-defined virtual devices (such as virtual device mappings) in case you have a system failure or disaster. In this situation, back up user-defined virtual devices by saving the data to a location that is automatically backed up when you use the **backupios** command to back up the VIOS.

Related information

ioslevel command

viosbr command

Scheduling backups of the Virtual I/O Server and user-defined virtual devices

You can schedule regular backups of the Virtual I/O Server (VIOS) and user-defined virtual devices to ensure that your backup copy accurately reflects the current configuration.

About this task

To ensure that your backup of the VIOS accurately reflects your current running VIOS, back up the VIOS and the user-defined virtual devices each time that the configuration changes. For example:

- Changing the VIOS, like installing a fix pack.

- Adding, deleting, or changing the external device configuration, like changing the SAN configuration.
- Adding, deleting, or changing resource allocations and assignments for the VIOS, like memory, processors, or virtual and physical devices.
- Adding, deleting, or changing user-defined virtual device configurations, like virtual device mappings.

You can schedule backups in one of the following ways:

- You can schedule backups of the VIOS and user-defined virtual devices by creating a script that includes the **backupios** command. Then, create a crontab file entry that runs the script on a regular interval. Use this option in situations where you plan to restore the VIOS to a new or different system. (For example, use this option in the event of a system failure or disaster.)
- You can schedule backups of the configuration information for the user-defined virtual devices by using the **viosbr** command. Use this option in situations where you plan to restore the configuration information to the same VIOS partition from which it was backed up.

Scheduling backups of the Virtual I/O Server and user-defined virtual devices by creating a script and crontab file entry

You can schedule regular backups of the Virtual I/O Server (VIOS) and user-defined virtual devices to ensure that your backup copy accurately reflects the current configuration.

About this task

To ensure that your backup of the VIOS accurately reflects your current running VIOS, back up the VIOS each time that its configuration changes. For example:

- Changing the VIOS, like installing a fix pack.
- Adding, deleting, or changing the external device configuration, like changing the SAN configuration.
- Adding, deleting, or changing resource allocations and assignments for the VIOS, like memory, processors, or virtual and physical devices.
- Adding, deleting, or changing user-defined virtual device configurations, like virtual device mappings.

Before you start, ensure that you are logged in to the VIOS as the prime administrator (padmin).

To back up the VIOS and user-defined virtual devices, complete the following tasks:

Procedure

1. Create a script for backing up the VIOS, and save it in a directory that is accessible to the **padmin** user ID. For example, create a script called *backup* and save it in the `/home/padmin` directory.

   Ensure that your script includes the following information:

   - The **backupios** command for backing up the VIOS.
   - Commands for saving information about user-defined virtual devices.
   - Commands to save the virtual devices information to a location that is automatically backed up when you use the **backupios** command to back up the VIOS.
2. Create a **crontab** file entry that runs the *backup* script on a regular interval. For example, to run *backup* every Saturday at 2:00 AM., type the following commands:

   a. `crontab -e`

   b. `0 2 * * 6 /home/padmin/backup`

   When you complete the task, remember to save and exit.

Related information

backupios command

crontab command

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Scheduling backups of user-defined virtual devices by using the `viosbr` command

You can schedule regular backups of the user-defined virtual devices on the Virtual I/O Server (VIOS) logical partition. Scheduling regular backups ensures that your backup copy accurately reflects the current configuration.

About this task

To ensure that your backup of the user-defined virtual devices accurately reflects your currently running VIOS, back up the configuration information of the user-defined virtual devices each time that the configuration changes.

Before you start, run the **`ioslevel`** command to verify that the VIOS is at Version 2.1.2.0, or later.

Procedure

To back up the configuration information of the user-defined virtual devices, run the **`viosbr`** command as follows:

```
viosbr -backup -file /tmp/myserverbackup -frequency how_often
```

where,

- */tmp/myserverbackup* is the file to which you want to back up the configuration information.
- *how_often* is the frequency with which you want to back up the configuration information. You can specify one of the following values:
  - `daily`: Daily backups occur every day at 00:00.
  - `weekly`: Weekly backups occur every Sunday at 00:00.
  - `monthly`: Monthly backups occur on the first day of every month at 00:01.

Related tasks

Backing up user-defined virtual devices by using the viosbr command

You can back up user-defined virtual devices by using the **`viosbr`** command. Use the **`viosbr`** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Related information

ioslevel command

viosbr command

Backing up the Virtual I/O Server by using IBM Tivoli Storage Manager

You can use the IBM Tivoli Storage Manager to automatically back up the Virtual I/O Server on regular intervals, or you can perform incremental backups.

Backing up the Virtual I/O Server by using IBM Tivoli Storage Manager automated backup

You can automate backups of the Virtual I/O Server by using the **`crontab`** command and the IBM Tivoli Storage Manager scheduler.

About this task

Before you start, complete the following tasks:

- Ensure that you configured the Tivoli Storage Manager client on the Virtual I/O Server. For instructions, see "Configuring the IBM Tivoli Storage Manager client" on page 189.
- Ensure that you are logged in to the Virtual I/O Server as the prime administrator (padmin).

To automate backups of the Virtual I/O Server, complete the following steps:

Procedure

1. Write a script that creates an mksysb image of the Virtual I/O Server and save it in a directory that is accessible to the **padmin** user ID. For example, create a script called *backup* and save it in the `/home/padmin` directory. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then ensure that your script includes commands for saving information about user-defined virtual devices.

   For more information, see the following tasks:

   - For instructions about how to create an mksysb image, see "Backing up the Virtual I/O Server to a remote file system by creating an mksysb image" on page 209.
   - For instructions about how to save user-defined virtual devices, see "Backing up user-defined virtual devices by using the backupios command" on page 210.
2. Create a **`crontab`** file entry that runs the *backup* script on a regular interval. For example, to create an mksysb image every Saturday at 2:00 AM, type the following commands:

   a. `crontab -e`

   b. `0 2 0 0 6 /home/padmin/backup`

   When you are finished, remember to save and exit.
3. Work with the Tivoli Storage Manager administrator to associate the Tivoli Storage Manager client node with one or more schedules that are part of the policy domain.

   This task is not performed on the Tivoli Storage Manager client on the Virtual I/O Server. This task is performed by the Tivoli Storage Manager administrator on the Tivoli Storage Manager server.
4. Start the client scheduler and connect to the server schedule by using the **dsmc** command as follows:

   ```
   dsmc -schedule
   ```
5. If you want the client scheduler to restart when the Virtual I/O Server restarts, then add the following entry to the `/etc/inittab` file:

   ```
   itsm::once:/usr/bin/dsmc sched > /dev/null 2>&1 # TSM scheduler
   ```

Related information

IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

Backing up the Virtual I/O Server by using IBM Tivoli Storage Manager incremental backup

You can back up the Virtual I/O Server at any time by performing an incremental backup with the IBM Tivoli Storage Manager.

About this task

Perform incremental backups in situations where the automated backup does not suit your needs. For example, before you upgrade the Virtual I/O Server, perform an incremental backup to ensure that you have a backup of the current configuration. Then, after you upgrade the Virtual I/O Server, perform another incremental backup to ensure that you have a backup of the upgraded configuration.

Before you start, complete the following tasks:

- Ensure that you configured the Tivoli Storage Manager client on the Virtual I/O Server. For instructions, see "Configuring the IBM Tivoli Storage Manager client" on page 189.
- Ensure that you have an mksysb image of the Virtual I/O Server. If you plan to restore the Virtual I/O Server to a different system from which it was backed up, then ensure that the mksysb includes information about user-defined virtual devices. For more information, see the following tasks:
  - For instructions about how to create an mksysb image, see "Backing up the Virtual I/O Server to a remote file system by creating an mksysb image" on page 209.

- For instructions about how to save user-defined virtual devices, see "Backing up user-defined virtual devices by using the backupios command" on page 210.

Procedure

To perform an incremental backup of the Virtual I/O Server, run the **dsmc** command.

For example,

```
dsmc -incremental sourcefilespec
```

Where *sourcefilespec* is the directory path to where the `mksysb` file is located. For example, `/home/padmin/mksysb_image`.

Related information

IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Installation and User's Guide

Restoring the Virtual I/O Server

You can restore the Virtual I/O Server (VIOS) and user-defined virtual devices by using the **installios** command, the **viosbr** command, or IBM Tivoli Storage Manager.

About this task

The VIOS contains the following types of information that you need to restore: the VIOS itself and user-defined virtual devices.

- The VIOS includes the base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata. All this information is restored when you use the **installios** command.
- User-defined virtual devices include metadata, such as virtual devices mappings, that define the relationship between the physical environment and the virtual environment. You can restore user-defined virtual devices in one of the following ways:
  - You can restore user-defined virtual devices by using the **viosbr** command. Use this option in situations where you plan to restore the configuration information to the same VIOS partition from which it was backed up.
  - You can restore user-defined virtual devices by restoring the volume groups and manually re-creating virtual device mappings. Use this option in situations where you plan to restore the VIOS to a new or different system. (For example, in the event of a system failure or disaster.) Furthermore, in these situations, you also need to restore the following components of your environment. Back up these components to fully recover your VIOS configuration:
    - External device configurations, such as storage area network (SAN) devices.
    - Resources that are defined on the Hardware Management Console (HMC), such as processor and memory allocations. In other words, restore your HMC partition profile data for the VIOS and its client partitions.
    - The operating systems and applications that run in the client logical partitions.

  Note: To perform Live Partition Mobility after you restore the VIOS, ensure that you restart the HMC.

You can back up and restore the VIOS as follows.

Table 47. Backup and restoration methods for the VIOS

Backup method	Media	Restoration method
To tape	Tape	From tape
To DVD	DVD-RAM	From DVD

Table 47. Backup and restoration methods for the VIOS (continued)

Backup method	Media	Restoration method
To remote file system	nim_resources.tar image	From an HMC using the Network Installation Management (NIM) on Linux facility and the **installios** command
To remote file system	mksysb image	From an AIX 5L NIM server, or later and a standard mksysb system installation
To remote file system	mksysb image	From an AIX 5L NIM server, or later and a standard mksysb system installation
Tivoli Storage Manager	mksysb image	Tivoli Storage Manager

What to do next

Related tasks

Backing up the Virtual I/O Server

You can back up the Virtual I/O Server (VIOS) and user-defined virtual devices by using the **backupios** command or the **viosbr** command. You can also use IBM Tivoli Storage Manager to schedule backups and to store backups on another server.

Related information

installios command

viosbr command

Restoring the Virtual I/O Server from tape

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from tape.

About this task

To restore the Virtual I/O Server from tape, follow these steps:

Procedure

1. Specify the Virtual I/O Server logical partition to boot from the tape by using the **bootlist** command. Alternatively, you can alter the bootlist in the System Management Services (SMS).
2. Insert the tape into the tape drive.
3. From the **SMS** menu, select to install from the tape drive.
4. Follow the installation steps according to the system prompts.
5. If you restored the Virtual I/O Server to a different system from which it was backed up, then you need to restore the user-defined virtual devices.

   For instructions, see "Restoring user-defined virtual devices manually" on page 221.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Restoring the Virtual I/O Server from one or more DVDs

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from one or more DVDs.

About this task

To restore the Virtual I/O Server from one or more DVDs, follow these steps:

Procedure

1. Specify the Virtual I/O Server partition to boot from the DVD by using the **bootlist** command. Alternatively, you can alter the bootlist in the System Management Services (SMS).
2. Insert the DVD into the optical drive.
3. From the **SMS** menu, select to install from the optical drive.
4. Follow the installation steps according to the system prompts.
5. If you restored the Virtual I/O Server to a different system from which it was backed up, then you need to restore the user-defined virtual devices.

   For instructions, see "Restoring user-defined virtual devices manually" on page 221.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Restoring the Virtual I/O Server from the HMC by using a `nim_resources.tar` file

You can restore the Virtual I/O Server (VIOS) base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from a `nim_resources.tar` image stored in a remote file system.

About this task

To restore the Virtual I/O Server from a `nim_resources.tar` image in a file system, complete the following steps:

Procedure

1. Run the **`installios`** command from the HMC command line.

   This restores a backup image, `nim_resources.tar`, that was created by using the **`backupios`** command.
2. Follow the installation procedures according to the system prompts. The source of the installation images is the exported directory from the backup procedure. For example, `server1:/export/ios_backup`.
3. When the restoration is finished, open a virtual terminal connection (for example, by using telnet) to the Virtual I/O Server that you restored. Some additional user input might be required.
4. If you restored the Virtual I/O Server to a different system from which it was backed up, you must restore the user-defined virtual devices.

   For instructions, see "Restoring user-defined virtual devices manually" on page 221.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Restoring the Virtual I/O Server from a NIM server by using an mksysb file

You can restore the Virtual I/O Server base code, applied fix packs, custom device drivers to support disk subsystems, and some user-defined metadata from an mksysb image stored in a remote file system.

Before you begin

Before you start, complete the following tasks:

- Ensure that the server to which you plan to restore the Virtual I/O Server is defined as a Network Installation Management (NIM) resource.
- Ensure that the mksysb file (that contains the backup of the Virtual I/O Server) is on the NIM server.

About this task

To restore the Virtual I/O Server from an mksysb image in a file system, complete the following tasks:

Procedure

1. Define the mksysb file as a NIM resource, specifically, a NIM object, by running the **nim** command.

   To view a detailed description of the **nim** command, see nim Command.

   For example:

   ```
   nim -o define -t mksysb -a server=servername -alocation=/export/ios_backup/
   filename.mksysb objectname
   ```

   Where:

   - *servername* is the name of the server that holds the NIM resource.
   - *filename* is the name of the mksysb file.
   - *objectname* is the name by which NIM registers and recognizes the mksysb file.
2. Define a Shared Product Object Tree (SPOT) resource for the mksysb file by running the **nim** command. For example:

   ```
   nim -o define -t spot -a server=servername -a location=/export/ios_backup/
   SPOT -a source=objectname SPOTname
   ```

   Where:

   - *servername* is the name of the server that holds the NIM resource.
   - *objectname* is the name by which NIM registers and recognizes the mksysb file.
   - *SPOTname* is the NIM object name for the mksysb image that was created in the previous step.
3. Install the Virtual I/O Server from the mksysb file by using the **smit** command. For example:

   ```
   smit nim_bosinst
   ```

   Ensure that the following entry fields contain the following specifications.

   Table 48. Specifications for the SMIT command

Field	Specification
Installation TYPE	mksysb
SPOT	*SPOTname* from step 3
MKSYSB	*objectname* from step 2
Remain NIM client after install?	no

4. Start the Virtual I/O Server logical partition.

For instructions, see step 3, Boot the Virtual I/O Server, of Installing the Virtual I/O Server using NIM.

5. If you restored the Virtual I/O Server to a different system from which it was backed up, you must restore the user-defined virtual devices.

   For instructions, see "Restoring user-defined virtual devices manually" on page 221.

Related information

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Using the NIM define operation

Defining a SPOT resource

Installing a client using NIM

Restoring user-defined virtual devices

You can restore user-defined virtual devices on the Virtual I/O Server (VIOS) by restoring volume groups and manually re-creating virtual device mappings. Alternatively, you can restore user-defined virtual devices by using the **viosbr** command.

About this task

User-defined virtual devices include metadata, such as virtual device mappings, that define the relationship between the physical environment and the virtual environment. You can restore user-defined virtual devices in one of the following ways:

- You can restore user-defined virtual devices by restoring volume groups and manually re-creating virtual device mappings. Use this option in situations where you plan to restore the VIOS to a new or different system. (For example, use this option in the event of a system failure or disaster.)
- You can restore user-defined virtual devices by using the **viosbr** command. Use this option in situations where you plan to restore the configuration information to the same VIOS partition from which it was backed up.

Related tasks

Backing up user-defined virtual devices
You can back up user-defined virtual devices by saving the data to a location that is automatically backed up when you use the **backupios** command to back up the Virtual I/O Server (VIOS). Alternatively, you can back up user-defined virtual devices by using the **viosbr** command.

Restoring user-defined virtual devices manually

In addition to restoring the Virtual I/O Server (VIOS), you might need to restore user-defined virtual devices (such as virtual device mappings). For example, in the event of a system failure, system migration, or disaster, you need to restore both the VIOS and user-defined virtual devices. In this situation, restore the volume groups by using the **restorevgstruct** command and manually re-create the virtual device mappings by using the **mkvdev** command.

Before you begin

User-defined virtual devices include metadata, such as virtual device mappings, that define the relationship between the physical environment and the virtual environment. In situations where you plan to restore the VIOS to a new or different system, you need to back up both the VIOS and user-defined virtual devices. (For example, in the event of a system failure or disaster, you must restore both the VIOS and user-defined virtual devices.)

Before you start, restore the VIOS from tape, DVD, or a remote file system. For instructions, see one of the following procedures:

- "Restoring the Virtual I/O Server from tape" on page 218
- "Restoring the Virtual I/O Server from one or more DVDs" on page 219
- "Restoring the Virtual I/O Server from the HMC by using a nim_resources.tar file" on page 219
- "Restoring the Virtual I/O Server from a NIM server by using an mksysb file" on page 220

About this task

To restore user-defined virtual devices, complete the following steps:

Procedure

1. List all the backed-up volume groups (or storage pools) by running the following command:

   ```
   restorevgstruct -ls
   ```

   This command lists the files that are located in the **/home/ios/vgbackups** directory.
2. Run the **lspv** command to determine which disks are empty.
3. Restore the volume groups (or storage pools) to the empty disks by running the following command for each volume group (or storage pool):

   ```
   restorevgstruct -vg volumegroup hdiskx
   ```

   Where:

   - *volumegroup* is the name of a volume group (or storage pool) from step 1.
   - *hdiskx* is the name of an empty disk from step 2.
4. Re-create the mappings between the virtual devices and physical devices by using the **mkvdev** command. Re-create mappings for storage device mappings, shared Ethernet and Ethernet adapter mappings, and virtual LAN settings. You can find mapping information in the file that you specified in the **tee** command from the backup procedure. For example, `/home/padmin/`*`filename`*.

Related tasks

Restoring user-defined virtual devices by using the viosbr command

You can restore user-defined virtual devices by using the **viosbr** command. Use the **viosbr** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Related information

mkvdev command

restorevgstruct command

tee command

IBM System p Advanced POWER Virtualization Best Practices RedPaper

Restoring user-defined virtual devices by using the `viosbr` command

You can restore user-defined virtual devices by using the **viosbr** command. Use the **viosbr** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Before you begin

The **viosbr** command restores the VIOS partition to the same state in which it was when the backup was taken. With the information available from the backup, the command performs the following actions:

- Sets the attribute values for physical devices, such as controllers, adapters, disks, optical devices, tape devices, and Ethernet interfaces.
- Imports logical devices, such as volume groups or storage pools, clusters, logical volumes, file systems, and repositories.
- Creates virtual devices and their corresponding mappings for devices like Etherchannel, Shared Ethernet Adapter, virtual target devices, virtual Fibre Channel adapters, and paging space devices.

Before you start, complete the following tasks:

1. Run the **ioslevel** command to verify that the VIOS is at Version 2.1.2.0, or later.
2. Determine the backup file that you want to restore. The backup file must be a file that was created by using the **viosbr -backup** command.

3. Verify that the VIOS partition to which you plan to restore the information is the same VIOS partition from which it was backed up.

Procedure

To restore all the possible devices and display a summary of deployed and nondeployed devices, run the following command:

```
viosbr -restore  -file /home/padmin/cfgbackups/myserverbackup.002.tar.gz
```

where, */home/padmin/cfgbackups/myserverbackup.002.tar.gz* is the backup file that contains the information that you want to restore.

The system displays information like the following output:

```
Backed up Devices that are unable to restore/change
==================================================
<Name(s) of non-deployed devices>
DEPLOYED or CHANGED devices:
============================
Dev name during BACKUP                  Dev name after RESTORE
-----------------------------           ------------------------------
<Name(s) of deployed devices>
```

Related tasks

Backing up user-defined virtual devices by using the viosbr command
You can back up user-defined virtual devices by using the **viosbr** command. Use the **viosbr** command when you plan to restore the information to the same Virtual I/O Server (VIOS) logical partition from which it was backed up.

Restoring user-defined virtual devices manually
In addition to restoring the Virtual I/O Server (VIOS), you might need to restore user-defined virtual devices (such as virtual device mappings). For example, in the event of a system failure, system migration, or disaster, you need to restore both the VIOS and user-defined virtual devices. In this situation, restore the volume groups by using the **restorevgstruct** command and manually re-create the virtual device mappings by using the **mkvdev** command.

Related information

ioslevel command

viosbr command

Restoring the Virtual I/O Server by using IBM Tivoli Storage Manager

You can use the IBM Tivoli Storage Manager to restore the mksysb image of the Virtual I/O Server.

About this task

You can restore the Virtual I/O Server to the system from which it was backed up, or to a new or different system (for example, in the event of a system failure or disaster). The following procedure applies to restoring the Virtual I/O Server to the system from which it was backed up. First, you restore the mksysb image to the Virtual I/O Server by using the **dsmc** command on the Tivoli Storage Manager client. But restoring the mksysb image does not restore the Virtual I/O Server. You then need to transfer the mksysb image to another system and convert the mksysb image to an installable format.

To restore the Virtual I/O Server to a new or different system, use one of the following procedures:

- "Restoring the Virtual I/O Server from tape" on page 218
- "Restoring the Virtual I/O Server from one or more DVDs" on page 219
- "Restoring the Virtual I/O Server from the HMC by using a nim_resources.tar file" on page 219
- "Restoring the Virtual I/O Server from a NIM server by using an mksysb file" on page 220

Before you start, complete the following tasks:

1. Ensure that the system to which you plan to transfer the mksysb image is running AIX.
2. Ensure that the system running AIX has a DVD-RW or CD-RW drive.
3. Ensure that AIX has the cdrecord and mkisofs RPMs downloaded and installed. To download and install the RPMs, see the AIX Toolbox for Linux Applications website.

Restriction: Interactive mode is not supported on the Virtual I/O Server. You can view session information by typing `dsmc` on the Virtual I/O Server command line.

To restore the Virtual I/O Server by using Tivoli Storage Manager, complete the following tasks:

Procedure

1. Determine which file you want to restore by running the **dsmc** command to display the files that have been backed up to the Tivoli Storage Manager server:

   ```
   dsmc -query
   ```

2. Restore the mksysb image by using the **dsmc** command.

   For example:

   ```
   dsmc -restore sourcefilespec
   ```

   Where *sourcefilespec* is the directory path to the location where you want to restore the mksysb image. For example, /home/padmin/mksysb_image

3. Transfer the mksysb image to a server with a DVD-RW or CD-RW drive by running the following File Transfer Protocol (FTP) commands:

   a) Run the following command to make sure that the FTP server is started on the Virtual I/O Server: `startnetsvc ftp`

   b) Open an FTP session to the server with the DVD-RW or CD-RW drive: `ftp` *`server_hostname`*, where *server_hostname* is the host name of the server with the DVD-RW or CD-RW drive.

   c) At the FTP prompt, change to the installation directory to the directory where you want to save the mksysb image.

   d) Set the transfer mode to binary: `binary`

   e) Turn off interactive prompting if it is on: `prompt`

   f) Transfer the mksysb image to the server: `mput mksysb_image`

   g) Close the FTP session, after transferring mksysb image, by typing `quit`.

4. Write the mksysb image to CD or DVD by using the **mkcd** or **mkdvd** commands.
5. Reinstall the Virtual I/O Server by using the CD or DVD that you created.

   For instructions, see "Restoring the Virtual I/O Server from one or more DVDs" on page 219.

Related reference

mkcd Command

mkdvd Command

Installing or replacing an adapter with the system power turned on in a Virtual I/O Server

Find information about how to install or replace an adapter in the Virtual I/O Server logical partition.

Before you begin

The Virtual I/O Server includes a Hot Plug Manager that is similar to the Hot Plug Manager in the AIX operating system. The Hot Plug Manager allows you to hot plug adapters into the system and then activate them for the logical partition without having to reboot the system. Use the Hot Plug Manager for adding, identifying, or replacing adapters in the system that are currently assigned to the Virtual I/O Server.

Prerequisites:

- If you are installing a new adapter, an empty system slot must be assigned to the Virtual I/O Server logical partition. This task can be done through dynamic logical partitioning (DLPAR) operations.
- If you are using a Hardware Management Console (HMC), you must also update the logical partition profile of the Virtual I/O Server so that the new adapter is configured to the Virtual I/O Server after you restart the system.
- If you are installing a new adapter, ensure that you have the software required to support the new adapter and determine whether there are any existing PTF prerequisites to install. For information about software prerequisites, see the IBM Prerequisite website (http://www-912.ibm.com/e_dir/eServerPrereq.nsf).

About this task

Choose from the following tasks:

- "Installing an adapter" on page 225
- "Replacing an adapter" on page 225
- "Unconfiguring storage adapters" on page 226
- "Preparing the client logical partitions" on page 227

Installing an adapter

About this task

To install an adapter with the system power on in Virtual I/O Server, complete the following steps:

Procedure

1. From the Hot Plug Manager, select **Add a PCIe Hot Plug Adapter**, then press Enter.

   The Add a Hot-Plug Adapter window is displayed.
2. Select the appropriate empty slot from those listed, and press Enter.

   A fast-blinking amber LED located at the back of the server near the adapter indicates that the slot has been identified.
3. Follow the instructions on the screen to install the adapter until the LED for the specified slot is set to the Action state.
   a. Set the adapter LED to the action state so that the indicator light for the adapter slot flashes
   b. Physically install the adapter
   c. Finish the adapter installation task in **diagmenu**.
4. Enter **cfgdev** to configure the device for the Virtual I/O Server.

Results

If you are installing a PCIe, Fibre Channel adapter, it is now ready to be attached to a SAN and have LUNs assigned to the Virtual I/O Server for virtualization.

Replacing an adapter

Before you begin

Prerequisite: Before you can remove or replace a storage adapter, you must unconfigure that adapter. See "Unconfiguring storage adapters" on page 226 for instructions.

About this task

To replace an adapter with the system power turned on in Virtual I/O Server, complete the following steps:

Procedure

1. From the PCIe Hot Plug Manager, select **Unconfigure a Device**, then press Enter.
2. Press F4 (or Esc +4) to display the **Device Names** menu.
3. Select the adapter you are removing in the **Device Names** menu.
4. In the **Keep Definition** field, use the Tab key to answer `Yes`. In the **Unconfigure Child Devices** field, use the Tab key again to answer YES, then press Enter.
5. Press Enter to verify the information on the **ARE YOU SURE** screen. Successful unconfiguration is indicated by the `OK` message displayed next to the Command field at the top of the screen.
6. Press F4 (or Esc +4) twice to return to the Hot Plug Manager.
7. Select **replace/remove PCIe Hot Plug adapter**.
8. Select the slot that has the device to be removed from the system.
9. Select **replace**.

   A fast-blinking amber LED located at the back of the machine near the adapter indicates that the slot has been identified.
10. Press Enter which places the adapter in the action state, meaning it is ready to be removed from the system.

Unconfiguring storage adapters

About this task

Before you can remove or replace a storage adapter, you must unconfigure that adapter. Storage adapters are generally parent devices to media devices, such as disk drives or tape drives. Removing the parent requires that all attached child devices either be removed or placed in the define state.

Unconfiguring a storage adapter involves the following tasks:

- Closing all applications that are using the adapter you are removing, replacing, or moving
- Unmounting file systems
- Ensuring that all devices connected to the adapter are identified and stopped
- Listing all slots that are currently in use or a slot that is occupied by a specific adapter
- Identifying the adapter's slot location
- Making parent and child devices unavailable
- Making the adapter unavailable

If the adapter supports physical volumes that are in use by a client logical partition, then You can perform steps on the client logical partition before unconfiguring the storage adapter. For instructions, see "Preparing the client logical partitions" on page 227. For example, the adapter might be in use because the physical volume was used to create a virtual target device, or it might be part of a volume group used to create a virtual target device.

To unconfigure SCSI, SSA, and Fibre Channel storage adapters, complete the following steps:

Procedure

1. Connect to the Virtual I/O Server command-line interface.
2. Enter `oem_setup_env` to close all applications that are using the adapter you are unconfiguring.
3. Type `lsslot-c PCI` to list all the hot plug slots in the system unit and display their characteristics.

4. Type `lsdev -C` to list the current state of all the devices in the system unit.
5. Type `unmount` to unmount previously mounted file systems, directories, or files using this adapter.
6. Type `rmdev -l adapter -R` to make the adapter unavailable.

Attention: Do not use the -d flag with the `rmdev` command for hot plug operations because this action removes your configuration.

Preparing the client logical partitions

About this task

If the virtual target devices of the client logical partitions are not available, the client logical partitions can fail or they might be unable to perform I/O operations for a particular application. If you use the HMC to manage the system, you might have redundant Virtual I/O Server logical partitions, which allow for Virtual I/O Server maintenance and avoid downtime for client logical partitions. If you are replacing an adapter on the Virtual I/O Server and your client logical partition is dependent on one or more of the physical volumes accessed by that adapter, then You can take action on the client before you unconfigure the adapter.

The virtual target devices must be in the define state before the Virtual I/O Server adapter can be replaced. Do not remove the virtual devices permanently.

Procedure

To prepare the client logical partitions so that you can unconfigure an adapter, complete the following steps depending on your situation.

Table 49. Situations and steps for preparing the client logical partitions

Situation	Steps
You have redundant hardware on the Virtual I/O Server for the adapter.	No action is required on the client logical partition.
HMC-managed systems only: You have redundant Virtual I/O Server logical partitions that, in conjunction with virtual client adapters, provide multiple paths to the physical volume on the client logical partition.	No action is required on the client logical partition. However, path errors might be logged on the client logical partition.
HMC-managed systems only: You have redundant Virtual I/O Server logical partitions that, in conjunction with virtual client adapters, provide multiple physical volumes that are used to mirror a volume group.	See the procedures for your client operating system. For example, for AIX, see Replacing a disk on the Virtual I/O Server in the Advanced POWER Virtualization Best Practices Redpaper. The procedure for Linux is similar to this procedure for AIX.For example, for AIX, see Replacing a disk on the Virtual I/O Server in the Advanced POWER Virtualization Best Practices Redpaper. The procedure for Linux is similar to this procedure for AIX.
You do not have redundant Virtual I/O Server logical partitions.	Shut down the client logical partition. For systems that are managed by the HMC, see Stopping a system (www.ibm.com/support/knowledgecenter/POWER9/p9haj/stopsyshmc.htm).

Viewing information and statistics about the Virtual I/O Server, the server, and virtual resources

You can view information and statistics about the Virtual I/O Server, the server, and virtual resources to help you manage and monitor the system, and troubleshoot problems.

About this task

The following table lists the information and statistics available on the Virtual I/O Server, as well as the commands you need to run to view the information and statistics.

Table 50. Information and associated commands for the Virtual I/O Server

Information to view	Command
Statistics about kernel threads, virtual memory, disks, traps, and processor activity.	**vmstat**
Statistics for a Fibre Channel device driver.	**fcstat**
A summary of virtual memory usage.	**svmon**
Information about the Virtual I/O Server and the server, such as the server model, machine ID, Virtual I/O Server logical partition name and ID, and the LAN network number.	**uname**
Generic and device-specific statistics for an Ethernet driver or device, including the following information for a Shared Ethernet Adapter: • Shared Ethernet Adapter statistics: – Number of real and virtual adapters (If you are using Shared Ethernet Adapter failover, this number does not include the control channel adapter) – Shared Ethernet Adapter flags – VLAN IDs – Information about real and virtual adapters • Shared Ethernet Adapter failover statistics: – High availability statistics – Packet types – State of the Shared Ethernet Adapter – Bridging mode • GARP VLAN Registration Protocol (GVRP) statistics: – Bridge Protocol Data Unit (BPDU) statistics – Generic Attribute Registration Protocol (GARP) statistics – GARP VLAN Registration Protocol (GVRP) statistics • Listing of the individual adapter statistics for the adapters that are associated with the Shared Ethernet Adapter	**enstat**

The **vmstat**, **fcstat**, **svmon**, and **uname** commands are available with Virtual I/O Server Version 1.5 or later. To update the Virtual I/O Server, see "Updating the Virtual I/O Server" on page 205.

Virtual I/O Server Performance Advisor

The VIOS Performance Advisor tool provides advisory reports that are based on the key performance metrics on various partition resources that are collected from the VIOS environment.

Starting with Virtual I/O Server (VIOS) Version 2.2.2.0, you can use the VIOS Performance Advisor tool. Use this tool to provide health reports that have proposals for making configurational changes to the VIOS environment and to identify areas to investigate further. On the VIOS command line, enter the `part` command to start the VIOS Performance Advisor tool.

You can start the VIOS Performance Advisor tool in the following modes:

- On-demand monitoring mode
- Postprocessing mode

When you start the VIOS Performance Advisor tool in the on-demand monitoring mode, provide the duration for which the tool must monitor the system in minutes. The duration that you provide must be 10 - 60 minutes at the end of which the tool generates the reports. During this time, samples are collected at regular intervals of 15 seconds. For example, to monitor the system for 30 minutes and generate a report, enter the following command:

```
part –i 30
```

Reports for the on-demand monitoring mode are successfully generated in the `ic43_120228_06_15_20.tar` file.

The output that is generated by the **part** command is saved in a .tar file, which is created in the current working directory. The naming convention for files in the on-demand monitoring mode is *short-hostname_yymmdd_hhmmss.tar*. In the postprocessing mode, the file name is that of the input file with the file name extension changed from an .nmon file to a .tar file.

When you start the VIOS Performance Advisor tool in the postprocessing mode, you must provide a file as the input. The tool tries to extract as much data as possible from the file that you provide, and the tool generates reports. If the file does not have the required data for the tool to generate reports, an `Insufficient Data` message is added to the relevant fields. For example, to generate a report based on the data available in the `ic43_120206_1511.`nmon file, enter the following command:

```
part -f ic43_120206_1511.nmon
```

Reports for the postprocessing mode are successfully generated in the `ic43_120206_1511.tar` file.

Note: The size of the input file in the postprocessing mode must be within 100 MB because postprocessing of huge data results in more time to generate the reports. For example, if the size of a file is 100 MB and the VIOS has 255 disks that are configured, with more than 4000 samples, it might take 2 minutes to generate the reports.

Related information

part command

Virtual I/O Server Performance Advisor reports

The Virtual I/O Server (VIOS) Performance Advisor tool provides advisory reports that are related to performance of various subsystems in the VIOS environment.

The output that is generated by the **part** command is saved in a .tar file that is created in the current working directory.

The `vios_advisor.xml` report is present in the output .tar file with the other supporting files. To view the generated report, complete the following steps:

1. Transfer the generated .tar file to a system that has a browser and a .tar file extractor installed.

2. Extract the .tar file.
3. Open the `vios_advisor.xml` file that is in the extracted directory.

The `vios_advisor.xml` file structure is based on an XML Schema Definition (XSD) in the `/usr/perf/analysis/vios_advisor.xsd` file.

Each report is shown in a tabular form, and the descriptions of all of the columns are provided in the following table.

Table 51. Performance metrics

Performance metrics	Description
Measured Value	This metric displays the values that are related to the performance metrics collected over a period.
Recommended Value	This metric displays all the suggested values when the performance metrics pass the critical thresholds.
First Observed	This metric displays the time stamp when the measured value is first observed.
Last Observed	This metric displays the time stamp when the measured value is last observed.
Risk	If either the warning or the critical thresholds are passed, the risk factor is indicated on a scale of 1 - 5 with 1 being the lowest value and 5 being the highest value.
Impact	If either the warning or critical thresholds are passed, the impact is indicated on a scale of 1 - 5 with 1 being the lowest value and 5 being the highest value.

The following are the types of advisory reports that are generated by the VIOS Performance Advisor tool:

- System configuration advisory report
- CPU (central processing unit) advisory report
- Memory advisory report
- Disk advisory report
- Disk adapter advisory report
- I/O activities (disk and network) advisory report

The system configuration advisory report consists of the information that is related to the VIOS configuration, such as processor family, server model, number of cores, frequency at which the cores are running, and the VIOS version. The output is similar to the following figure:

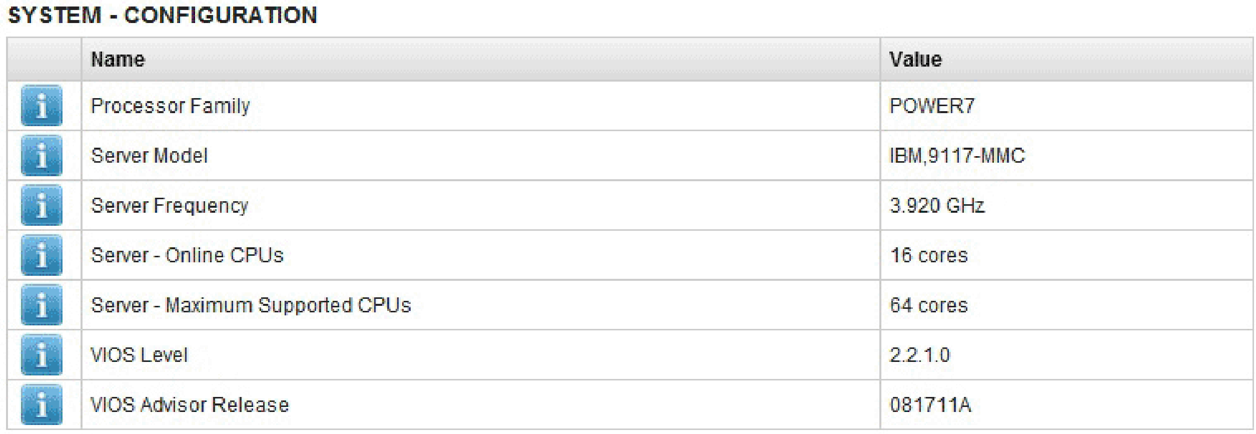

SYSTEM - CONFIGURATION

	Name	Value
ℹ	Processor Family	POWER7
ℹ	Server Model	IBM,9117-MMC
ℹ	Server Frequency	3.920 GHz
ℹ	Server - Online CPUs	16 cores
ℹ	Server - Maximum Supported CPUs	64 cores
ℹ	VIOS Level	2.2.1.0
ℹ	VIOS Advisor Release	081711A

The CPU advisory report consists of the information that is related to the processor resources, such as the number of cores assigned to the VIOS, processor consumption during the monitoring interval, and shared processor pool capacity for shared partitions. The output is similar to the following figure:

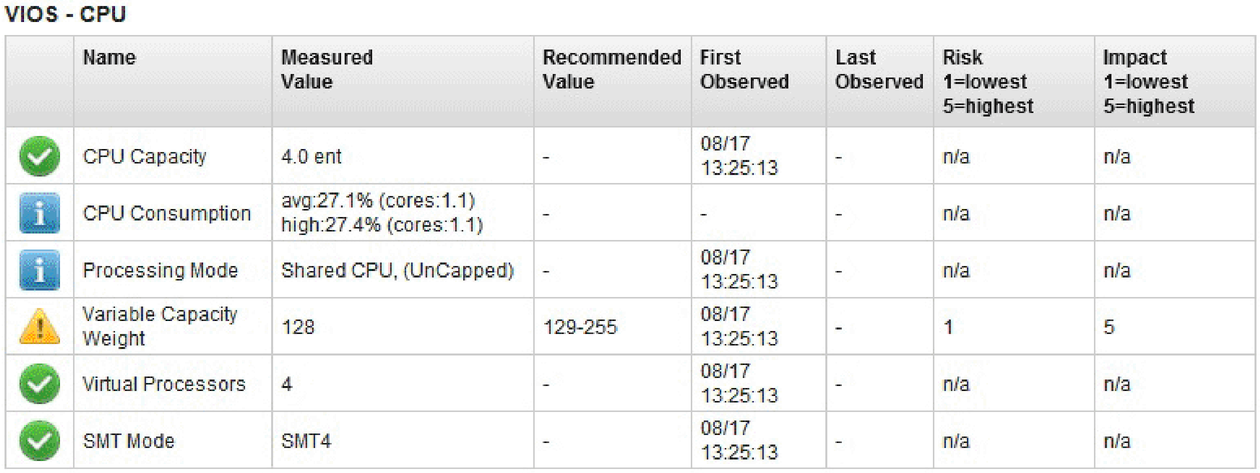

VIOS - CPU

	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
✓	CPU Capacity	4.0 ent	-	08/17 13:25:13	-	n/a	n/a
ℹ	CPU Consumption	avg:27.1% (cores:1.1) high:27.4% (cores:1.1)	-	-	-	n/a	n/a
ℹ	Processing Mode	Shared CPU, (UnCapped)	-	08/17 13:25:13	-	n/a	n/a
⚠	Variable Capacity Weight	128	129-255	08/17 13:25:13	-	1	5
✓	Virtual Processors	4	-	08/17 13:25:13	-	n/a	n/a
✓	SMT Mode	SMT4	-	08/17 13:25:13	-	n/a	n/a

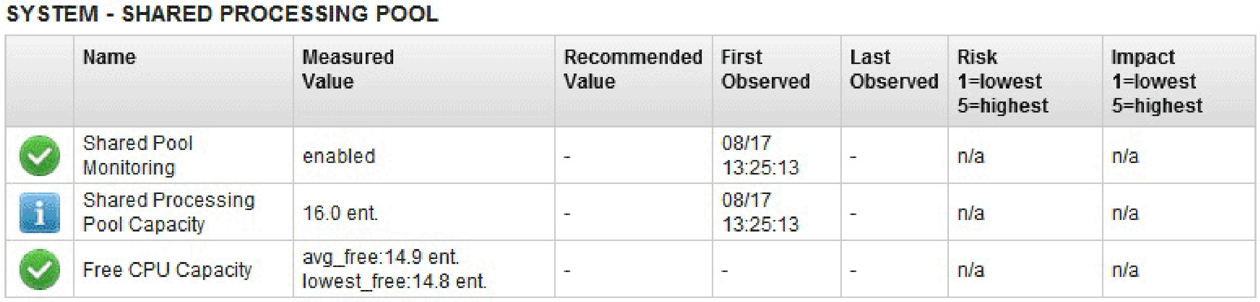

SYSTEM - SHARED PROCESSING POOL

	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
✓	Shared Pool Monitoring	enabled	-	08/17 13:25:13	-	n/a	n/a
ℹ	Shared Processing Pool Capacity	16.0 ent.	-	08/17 13:25:13	-	n/a	n/a
✓	Free CPU Capacity	avg_free:14.9 ent. lowest_free:14.8 ent.	-	-	-	n/a	n/a

Note: In the VIOS - CPU table, the status of the variable capacity weight is marked with the **Warning** icon because the best practice is for the VIOS to have an increased priority of 129 - 255 when in uncapped shared processor mode. See Table 52 on page 233 for the definitions about the **Warning** icon.

The memory advisory report consists of the information that is related to the memory resources, such as the available free memory, paging space that is allocated, paging rate, and pinned memory. The output is similar to the following figure:

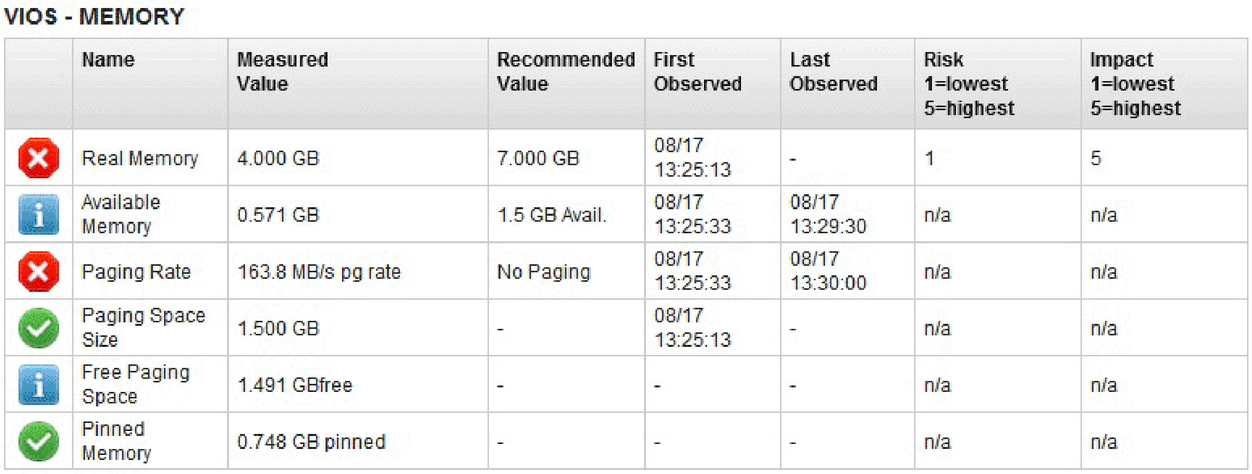

VIOS - MEMORY

	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
✖	Real Memory	4.000 GB	7.000 GB	08/17 13:25:13	-	1	5
ℹ	Available Memory	0.571 GB	1.5 GB Avail.	08/17 13:25:33	08/17 13:29:30	n/a	n/a
✖	Paging Rate	163.8 MB/s pg rate	No Paging	08/17 13:25:33	08/17 13:30:00	n/a	n/a
✓	Paging Space Size	1.500 GB	-	08/17 13:25:13	-	n/a	n/a
ℹ	Free Paging Space	1.491 GBfree	-	-	-	n/a	n/a
✓	Pinned Memory	0.748 GB pinned	-	-	-	n/a	n/a

Note: In this report, the status of the real memory is marked with the **Critical** icon because the available memory is less than the 1.5 GB limit that is specified in the Recommended Value column of the available memory. See Table 52 on page 233 for the definitions about the **Critical** icon.

The disk advisory report consists of the information that is related to the disks attached to the VIOS, such as the I/O activities that are getting blocked and I/O latencies. The output is similar to the following figure:

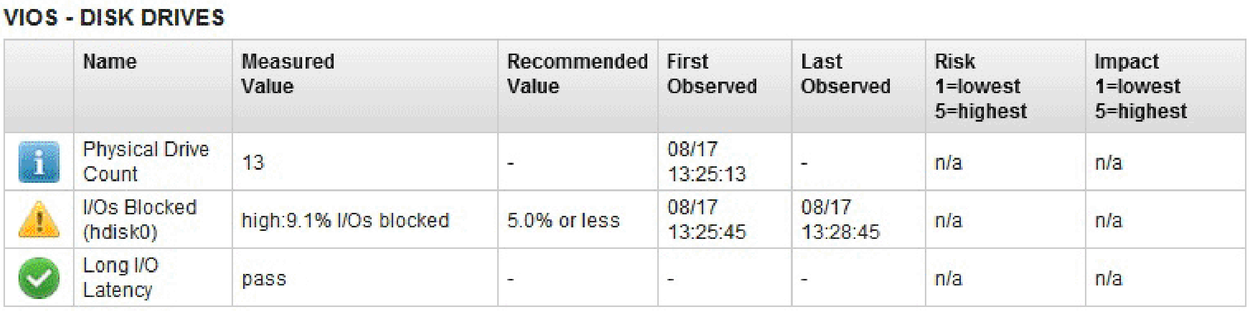

VIOS - DISK DRIVES

	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
ℹ	Physical Drive Count	13	-	08/17 13:25:13	-	n/a	n/a
⚠	I/Os Blocked (hdisk0)	high:9.1% I/Os blocked	5.0% or less	08/17 13:25:45	08/17 13:28:45	n/a	n/a
✓	Long I/O Latency	pass	-	-	-	n/a	n/a

The disk adapter advisory report consists of information that is related to the Fibre Channel adapters that are connected to the VIOS. This report illustrates the information that is based on the average I/O operations per second, adapter utilization, and running speed. The output is similar to the following figure:

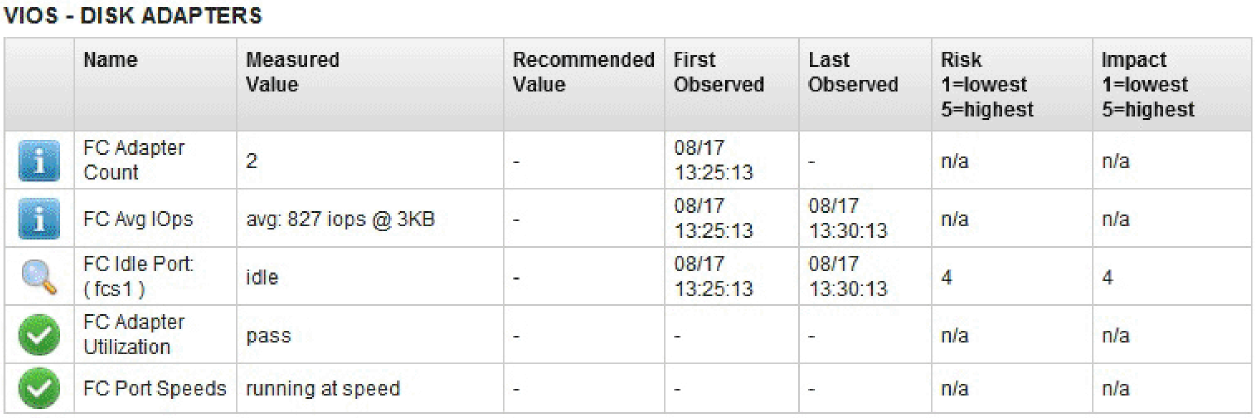

VIOS - DISK ADAPTERS

	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
ℹ	FC Adapter Count	2	-	08/17 13:25:13	-	n/a	n/a
ℹ	FC Avg IOps	avg: 827 iops @ 3KB	-	08/17 13:25:13	08/17 13:30:13	n/a	n/a
🔍	FC Idle Port (fcs1)	idle	-	08/17 13:25:13	08/17 13:30:13	4	4
✓	FC Adapter Utilization	pass	-	-	-	n/a	n/a
✓	FC Port Speeds	running at speed	-	-	-	n/a	n/a

Note: In this report, the status of the Fibre Channel idle port is marked with the **Investigate** icon because the tool identifies a Fibre Channel adapter that is not used often. See Table 52 on page 233 for the definitions about the **Investigate** icon.

The I/O activity advisory report consists of the following information:

- Disk I/O activity, such as average and peak I/O operations per second

- Network I/O activity, such as average and peak inflow and outflow I/O per second

The output is similar to the following figure:

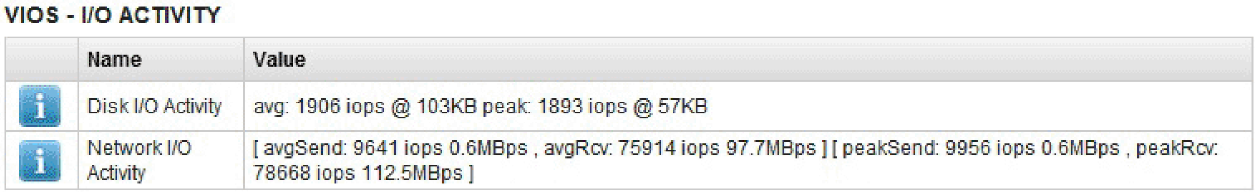

VIOS - I/O ACTIVITY

	Name	Value
ℹ	Disk I/O Activity	avg: 1906 iops @ 103KB peak: 1893 iops @ 57KB
ℹ	Network I/O Activity	[avgSend: 9641 iops 0.6MBps , avgRcv: 75914 iops 97.7MBps] [peakSend: 9956 iops 0.6MBps , peakRcv: 78668 iops 112.5MBps]

The details that are related to these advisory reports can also be obtained by clicking the respective report fields from the browser. The following details are available for all the advisory reports:

- What Is This: Brief description of the advisory field
- Why Important: Significance of the particular advisory field
- How to Modify: Details that are related to the configuration steps that you can use to modify the parameters that are related to the particular advisory field

For example, to know more about the processor capacity, you can click the corresponding row in the VIOS - CPU table and the information is displayed.

Note: The suggested values are based on the behavior during the monitoring period; therefore, the values can be used only as a guideline.

The following table describes the icon definitions.

Table 52. Icon definitions

Icons	**Definitions**
ℹ	Information that is related to configuration parameters
✓	Values acceptable in most cases
⚠	Possible performance problem
✖	Severe performance problem
🔍	Investigation required

Related information

part command

Monitoring the Virtual I/O Server

You can monitor the Virtual I/O Server by using error logs or IBM Tivoli Monitoring.

Error logs

AIX and Linux client logical partitions log errors against failing I/O operations. Hardware errors on the client logical partitions that are associated with virtual devices usually have corresponding errors that are logged on the server. However, if the failure is within the client logical partition, there will not be errors on the server.

Note: On Linux client logical partitions, if the algorithm for retrying Small Computer Serial Interface (SCSI) temporary errors is different from the algorithm that is used by AIX, the errors might not be recorded on the server.

IBM Tivoli Monitoring

With Virtual I/O Server V1.3.0.1 (fix pack 8.1), you can install and configure the IBM Tivoli Monitoring System Edition for System p agent on the Virtual I/O Server. With Tivoli Monitoring System Edition for IBM Power Systems, you can monitor the health and availability of multiple Power Systems servers (including the Virtual I/O Server) from the Tivoli Enterprise Portal. Tivoli Monitoring System Edition for Power Systems gathers data from the Virtual I/O Server, including data about physical volumes, logical volumes, storage pools, storage mappings, network mappings, real memory, processor resources, mounted file system sizes, and so on. From the Tivoli Enterprise Portal, you can view a graphical representation of the data, use predefined thresholds to alert you on key metrics, and resolve issues based on suggestions that are provided by the Expert Advice feature of Tivoli Monitoring.

Security on the Virtual I/O Server

Become familiar with the Virtual I/O Server security features.

Beginning with Version 1.3 of the Virtual I/O Server, you can set security options that provide tighter security controls over your Virtual I/O Server environment. These options allow you to select a level of system security hardening and specify the settings allowable within that level. The Virtual I/O Server security feature also allows you to control network traffic by enabling the Virtual I/O Server firewall. You can configure these options by using the **viosecure** command. To help you set up system security when you initially install the Virtual I/O Server, the Virtual I/O Server provides the configuration assistance menu. You can access the configuration assistance menu by running the **cfgassist** command.

Using the **viosecure** command, you can set, change, and view current security settings. By default, no Virtual I/O Server security levels are set. You must run the **viosecure** command to change the settings.

The following sections provide an overview of these features.

Virtual I/O Server system security hardening

The system security hardening feature protects all elements of a system by tightening security or implementing a higher level of security. Although hundreds of security configurations are possible with the Virtual I/O Server security settings, you can easily implement security controls by specifying a high, medium, or low security level.

Using the system security hardening features provided by Virtual I/O Server, you can specify values such as the following:

- Password policy settings
- Actions such as usrck, pwdck, grpck, and sysck
- Default file-creation settings
- Settings included in the **crontab** command

Configuring a system at too high a security level might deny services that are needed. For example, telnet and rlogin are disabled for high-level security because the login password is sent over the network unencrypted. If a system is configured at too low a security level, the system might be vulnerable to security threats. Since each enterprise has its own unique set of security requirements, the predefined High, Medium, and Low security configuration settings are best suited as a starting point for security configuration rather than an exact match for the security requirements of a particular enterprise. As you become more familiar with the security settings, you can make adjustments by choosing the hardening rules that you want to apply. You can get information about the hardening rules by running the **man** command.

Virtual I/O Server firewall

Using the Virtual I/O Server firewall, you can enforce limitations on IP activity in your virtual environment. With this feature, you can specify which ports and network services are allowed access to the Virtual I/O Server system. For example, if you need to restrict login activity from an unauthorized port, you can specify the port name or number and specify deny removing it from the allow list. You can also restrict a specific IP address.

Connecting to the Virtual I/O Server by using OpenSSH

You can set up remote connections to the Virtual I/O Server by using secure connections.

About this task

You can use the Open Source Secure Sockets Layer (OpenSSL) and Portable Secure Shell (OpenSSH) software to connect to the Virtual I/O Server by using secure connections. For more information about OpenSSL and OpenSSH, see the OpenSSL Project and Portable SSH websites.

To connect to the Virtual I/O Server by using OpenSSH, complete the following tasks:

Procedure

1. If you are using a version of Virtual I/O Server before Version 1.3.0, then install OpenSSH before you connect. For instructions, see "Downloading, installing, and updating OpenSSH and OpenSSL" on page 236.
2. Connect to the Virtual I/O Server.

   If you are using Version 1.3.0 or later, then connect by using either an interactive or noninteractive shell. If you are using a version before 1.3.0, then connect by using only an interactive shell.

   - To connect by using an interactive shell, type the following command from the command line of a remote system:

     ```
     ssh username@vioshostname
     ```

     where, *username* is your user name for the Virtual I/O Server and *vioshostname* is the name of the Virtual I/O Server.
   - To connect by using a noninteractive shell, run the following command:

     ```
     ssh username@vioshostname command
     ```

     Where:

     – *username* is your user name for the Virtual I/O Server.

     – *vioshostname* is the name of the Virtual I/O Server.

     – *command* is the command that you want to run. For example, `ioscli lsmap -all`.

     Note: When you use a noninteractive shell, remember to use the full command form (including the `ioscli` prefix) for all Virtual I/O Server commands.
3. Authenticate SSH.

   If you are using Version 1.3.0 or later, then authenticate by using either passwords or keys. If you are using a version before 1.3.0, then authenticate by using only passwords.

   - To authenticate by using passwords, enter your user name and password when prompted by the SSH client.
   - To authenticate by using keys, perform the following steps on the SSH client's operating system:

     a. Create a directory called `$HOME/.ssh` to store the keys. You can use RSA or DSA keys.

     b. Run the **ssh-keygen** command to generate public and private keys. For example,

     ```
     ssh-keygen -t  rsa
     ```

     This creates the following files in the `$HOME/.ssh` directory:

     – Private key: id_rsa

     – Public key: id_rsa.pub

c. Run the following command to append the public key to the `authorized_keys2` file on the Virtual I/O Server:

```
cat $HOME/.ssh/public_key_file | ssh username@vioshostname tee -a /home/username/.ssh/
authorized_keys2
```

Where:

– *public_key_file* is the public key file that is generated in the previous step. For example, id_rsa.pub.

– *username* is your user name for the Virtual I/O Server.

– *vioshostname* is the name of the Virtual I/O Server.

What to do next

The Virtual I/O Server might not include the latest version of OpenSSH or OpenSSL with each release. In addition, there might be OpenSSH or OpenSSL updates released in between Virtual I/O Server releases. In these situations, you can update OpenSSH and OpenSSL on the Virtual I/O Server by downloading and installing OpenSSH and OpenSSL. For instructions, see "Downloading, installing, and updating OpenSSH and OpenSSL" on page 236.

Downloading, installing, and updating OpenSSH and OpenSSL

If you are using a Virtual I/O Server version before 1.3, you must download and install OpenSSH and OpenSSL software before you can connect to the Virtual I/O Server by using OpenSSH. You can also use this procedure to update OpenSSH and OpenSSL on the Virtual I/O Server.

About this task

OpenSSH and OpenSSL might need to be updated on your Virtual I/O Server if the Virtual I/O Server did not include the latest version of OpenSSH or OpenSSL, or if there were OpenSSH or OpenSSL updates released in between Virtual I/O Server releases. In these situations, you can update OpenSSH and OpenSSL on the Virtual I/O Server by downloading and installing OpenSSH and OpenSSL by using the following procedure.

For more information about OpenSSL and OpenSSH, see the OpenSSL Project and Portable SSH websites.

Downloading the Open Source software

About this task

The OpenSSL software contains the encrypted library that is required to use the OpenSSH software. To download the software, complete the following tasks:

Procedure

1. Download the OpenSSL package to your workstation or host computer.
   a) To get the package, go to the AIX Web Download Pack Programs website.
   b) If you are registered to download the packages, sign in and accept the license agreement.
   c) If you are not registered to download the packages, complete the registration process and accept the license agreement. After registering, you are redirected to the download page.
   d) Select the package for download: **openSSL** and click **Continue**.
   e) Select any version of the package and click **Download** now.
2. Download the OpenSSH software by completing the following steps:

   Note: Alternatively, you can install the software from the AIX Expansion Pack.

   a) From your workstation (or host computer), go to the SourceFORGE.net website.
   b) Click **Download OpenSSH on AIX** to view the latest file releases.

c) Select the appropriate download package and click **Download**.

d) Click the OpenSSH package (`tar.Z` file) to continue with the download.

3. Create a directory on the Virtual I/O Server for the Open Source software files.
   For example, to create an installation directory named install_ssh, run the following command: `mkdir install_ssh`.
4. Transfer the software packages to the Virtual I/O Server by running the following File Transfer Protocol (FTP) commands from the computer on which you downloaded the software packages:

   a) Run the following command to make sure that the FTP server is started on the Virtual I/O Server: `startnetsvc ftp`

   b) Open an FTP session to the Virtual I/O Server on your local host: `ftp` *`vios_server_hostname`*, where *vios_server_hostname* is the host name of the Virtual I/O Server.

   c) At the FTP prompt, change to the installation directory that you created for the Open Source files: `cd` *`install_ssh`*, where *install_ssh* is the directory that contains the Open Source files.

   d) Set the transfer mode to binary: `binary`

   e) Turn off interactive prompting if it is on: `prompt`

   f) Transfer the downloaded software to the Virtual I/O Server: `mput` *`ssl_software_pkg`*, where *ssl_software_pkg* is the software that you downloaded.

   g) Close the FTP session, after transferring both software packages, by typing `quit`.

Install the Open Source software on the Virtual I/O Server

About this task

To install the software, complete the following steps:

Procedure

1. Run the following command from the Virtual I/O Server command line: `updateios -dev` *`install_ssh`* `-accept -install`, where *install_ssh* is the directory that contains the Open Source files.

   The installation program automatically starts the Secure Shell daemon (sshd) on the server.
2. Begin by using the **ssh** and **scp** commands; no further configuration is required.

   Restrictions:

   - The **sftp** command is not supported on versions of Virtual I/O Server earlier than 1.3.
   - Noninteractive shells are not supported by using OpenSSH with the Virtual I/O Server versions, earlier than 1.3.

Configuring Virtual I/O Server system security hardening

Set the security level to specify security hardening rules for your Virtual I/O Server system.

Before you begin

To implement system security hardening rules, you can use the **`viosecure`** command to specify a security level of high, medium, or low. A default set of rules is defined for each level. You can also set a level of default, which returns the system to the system standard settings and removes any level settings that have been applied.

About this task

The low-level security settings are a subset of the medium level security settings, which are a subset of the high-level security settings. Therefore, the *high* level is the most restrictive and provides the greatest level of control. You can apply all of the rules for a specified level or select which rules to activate for your environment. By default, no Virtual I/O Server security levels are set; you must run the **viosecure** command to modify the settings.

Use the following tasks to configure the system security settings.

Setting a security level

Procedure

To set a Virtual I/O Server security level of high, medium, or low, use the command `viosecure -level`. For example:

```
viosecure -level low -apply
```

Changing the settings in a security level

Procedure

To set a Virtual I/O Server security level in which you specify which hardening rules to apply for the setting, run the **viosecure** command interactively.
For example:

a. At the Virtual I/O Server command line, type `viosecure -level high`. All the security level options (hardening rules) at that level are displayed ten at a time (pressing Enter displays the next set in the sequence).
b. Review the options that are displayed and make your selection by entering the numbers, which are separated by a comma, that you want to apply, or type **ALL** to apply all the options or **NONE** to apply none of the options.
c. Press **Enter** to display the next set of options, and continue entering your selections.

   Note: To exit the command without making any changes, type "q".

Viewing the current security setting

Procedure

To display the current Virtual I/O Server security level setting, use the **viosecure** command with the **-view** flag.
For example:

```
viosecure -view
```

Removing security level settings

Procedure

- To unset any previously set system security levels and return the system to the standard system settings, run the following command: `viosecure -level default`
- To remove the security settings that have been applied, run the following command: `viosecure -undo`

Configuring Virtual I/O Server firewall settings

Enable the Virtual I/O Server firewall to control IP activity.

Before you begin

The Virtual I/O Server firewall is not enabled by default. To enable the Virtual I/O Server firewall, you must turn it on by using the **viosecure** command with the **-firewall** option. When you enable it, the default setting is activated, which allows access for the following IP services:

- ftp
- ftp-data
- ssh
- web
- https
- rmc
- cimom

Note: The firewall settings are contained in the `viosecure.ctl` file in the `/home/ios/security` directory. If for some reason the `viosecure.ctl` file does not exist when you run the command to enable the firewall, you receive an error. You can use the **-force** option to enable the standard firewall default ports.

You can use the default setting or configure the firewall settings to meet the needs of your environment by specifying which ports or port services to allow. You can also turn off the firewall to deactivate the settings.

About this task

Use the following tasks at the Virtual I/O Server command line to configure the Virtual I/O Server firewall settings:

Procedure

1. Enable the Virtual I/O Server firewall by running the following command:

   ```
   viosecure -firewall on
   ```

2. Specify the ports to allow or deny, by using the following command:

   ```
   viosecure -firwall allow | deny -port number
   ```

3. View the current firewall settings by running the following command:

   ```
   viosecure -firewall view
   ```

4. If you want to disable the firewall configuration, run the following command:

   ```
   viosecure -firewall off
   ```

Configuring a Kerberos client on the Virtual I/O Server

You can configure a Kerberos client on the Virtual I/O Server to enhance security in communications across the internet.

Before you begin

Before you start, ensure that the Virtual I/O Server is at Version 1.5, or later. To update the Virtual I/O Server, see “Updating the Virtual I/O Server” on page 205.

About this task

Kerberos is a network authentication protocol that provides authentication for client and server applications by using a secret-key cryptography. It negotiates authenticated, and optionally encrypted, communications between two points anywhere on the internet. Kerberos authentication generally works as follows:

1. A Kerberos client sends a request for a ticket to the Key Distribution Center (KDC).
2. The KDC creates a ticket-granting ticket (TGT) for the client and encrypts it using the client's password as the key.
3. The KDC returns the encrypted TGT to the client.
4. The client attempts to decrypt the TGT by using its password.
5. If the client successfully decrypts the TGT (for example, if the client gives the correct password), the client keeps the decrypted TGT. The TGT indicates proof of the client's identity.

Procedure

To configure a Kerberos client on the Virtual I/O Server, run the following command.

```
mkkrb5clnt -c KDC_server -r realm_name \ -s Kerberos_server -d Kerberos_client
```

Where:

- *KDC_server* is the name of the KDC server.
- *realm_name* is the name of the realm to which you want to configure the Kerberos client.
- *Kerberos_server* is the fully qualified host name of the Kerberos server.
- *Kerberos_client* is the domain name of the Kerberos client.

For example:

```
mkkrb5clnt -c bob.kerberso.com -r KERBER.COM \ -s bob.kerberso.com -d testbox.com
```

In this example, you configure the Kerberos client, testbox.com, to the Kerberos server, bob.kerberso.com. The KDC is running on bob.kerberso.com.

Using role-based access control with the Virtual I/O Server

With Virtual I/O Server Version 2.2, and later, a system administrator can define roles based on job functions in an organization by using role-based access control (RBAC).

A system administrator can use role-based access control (RBAC) to define roles for users in the Virtual I/O Server. A role confers a set of permissions or authorizations to the assigned user. Thus, a user can perform only a specific set of system functions depending on the access rights that are given. For example, if the system administrator creates the role **UserManagement** with authorization to access user management commands and assigns this role to a user, that user can manage users on the system but has no further access rights.

The benefits of using role-based access control with the Virtual I/O Server are as follows:

- Splitting system management functions
- Providing better security by granting only necessary access rights to users
- Implementing and enforcing system management and access control consistently
- Managing and auditing system functions with ease

Authorizations

The Virtual I/O Server creates authorizations that closely emulate the authorizations of the AIX operating system. The Virtual I/O Server creates authorizations that closely emulate the authorizations of the AIX operating system. The authorizations emulate naming conventions and descriptions, but are only applicable to the Virtual I/O Server specific requirements. By default, the **padmin** user is granted all the authorizations on the Virtual I/O Server, and can run all the commands. The other types of users (created by using the **mkuser** command) retain their command execution permissions.

The **mkauth** command creates a new user-defined authorization in the authorization database. You can create authorization hierarchies by using a dot (.) in the *auth* parameter to create an authorization of the form *ParentAuth.SubParentAuth.SubSubParentAuth....* All parent elements in the *auth* parameter must exist in the authorization database before the authorization is created. The maximum number of parent elements that you can use to create an authorization is eight.

You can set authorization attributes when you create authorizations through the *Attribute=Value* parameter. Every authorization that you create must have a value for the **id** authorization attribute. If you do not specify the **id** attribute by using the **mkauth** command, the command automatically generates a unique ID for the authorization. If you specify an ID, the value must be unique and greater than 15000. The IDs 1 - 15000 are reserved for system-defined authorizations.

Naming convention:

The system-defined authorizations in the Virtual I/O Server start with **vios.**. Hence, user-defined authorizations must not start with **vios.** or **aix.** or **aix.**. Since the authorizations that start with **vios.** and **aix.** and **aix.** are considered system-defined authorizations, users cannot add any further hierarchies to these authorizations.

Restriction:

Unlike in the AIX operating system, users cannot create authorizations for all Virtual I/O Server commands. In the AIX operating system, an authorized user can create a hierarchy of authorizations for all the commands. However, in the Virtual I/O Server, authorizations can be created only for the commands or scripts that are owned by the user. Users cannot create any authorizations that start with **vios.** or **aix.** since they are considered system-defined authorizations. Hence, users cannot add any further hierarchies to these authorizations.

Users cannot create authorizations for all Virtual I/O Server commands. In the Virtual I/O Server, authorizations can be created only for the commands or scripts that are owned by the user. Users cannot create any authorizations that start with **vios.** since they are considered system-defined authorizations. Hence, users cannot add any further hierarchies to these authorizations.

Authorization names must not begin with a dash (-), plus sign (+), at sign (@), or tilde (~). They must not contain spaces, tabs, or newline characters. You cannot use the keywords **ALL**, **default**, **ALLOW_OWNER**, **ALLOW_GROUP**, **ALLOW_ALL**, or an asterisk (*) as an authorization name. Do not use the following characters within an authorization string:

- : (colon)
- " (quotation mark)
- # (number sign)
- , (comma)
- = (equal sign)
- \ (backslash)
- / (forward slash)
- ? (question mark)
- ' (single quotation mark)
- ` (grave accent)

The following table lists the authorizations corresponding to the Virtual I/O Server commands. The vios and subsequent child authorizations, for example, vios and vios.device are not used. If a user is given a role that has either the parent or subsequent child authorization, for example, vios or vios.device, that user will have access to all the subsequent children authorizations and their related commands. For example, a role with the authorization vios.device, gives the user access to all vios.device.config and vios.device.manage authorizations and their related commands.

Table 53. Authorizations corresponding to Virtual I/O Server commands

Command	Command options	Authorization
activatevg	All	vios.lvm.manage.varyon
alert	All	vios.system.cluster.alert
alt_root_vg	All	vios.lvm.change.altrootvg
artexdiff	All	vios.system.rtexpert.diff
artexget	All	vios.system.rtexpert.get
artexlist	All	vios.system.rtexpert.list
artexmerge	All	vios.system.rtexpert.merge
artexset	All	vios.system.rtexpert.set
backup	All	vios.fs.backup
backupios	All	vios.install.backup
bootlist	All	vios.install.bootlist
cattracerpt	All	vios.system.trace.format
cfgassist	All	vios.security.cfgassist
cfgdev	All	vios.device.config
cfglnagg	All	vios.network.config.lnagg
cfgnamesrv	All	vios.system.dns
cfgsvc	All	vios.system.config.agent
chauth	All	vios.security.auth.change
chbdsp	All	vios.device.manage.backing.change
chdate	All	vios.system.config.date.change
chdev	All	vios.device.manage.change
checkfs	All	vios.fs.check
chedition	All	vios.system.edition
chkdev	All	vios.device.manage.check
chlang	All	vios.system.config.locale
chlv	All	vios.lvm.manage.change
chpath	All	vios.device.manage.path.change
chrep	All	vios.device.manage.repos.change
chrole	All	vios.security.role.change
chsp	All	vios.device.manage.spool.change
chtcpip	All	vios.network.tcpip.change
chuser	All	vios.security.user.change
chvg	All	vios.lvm.manage.change
chvlog	All	vios.device.manage.vlog.change

Table 53. Authorizations corresponding to Virtual I/O Server commands (continued)

Command	Command options	Authorization
chvlrepo	All	vios.device.manage.vlrepo.change
chvopt	All	vios.device.manage.optical.change
cl_snmp	All	vios.security.manage.snmp.query
cleandisk	All	vios.system.cluster.change
cluster	All	vios.system.cluster.create
cplv	All	vios.lvm.manage.copy
cpvdi	All	vios.lvm.manage.copy
deactivatevg	All	vios.lvm.manage.varyoff
diagmenu	All	vios.system.diagnostics
dsmc	All	vios.system.manage.tsm
entstat	All	vios.network.stat.ent
errlog	**-rm**	vios.system.log
	Others	vios.system.log.view
exportvg	All	vios.lvm.manage.export
extendlv	All	vios.lvm.manage.extend
extendvg	All	vios.lvm.manage.extend
failgrp	**-create**, **-modify**, **-remove**	vios.device.manage.spool.change or vios.system.cluster.pool.modify
fcstat	All	vios.network.stat.fc
fsck	All	vios.fs.check
hostmap	All	vios.system.config.address
hostname	All	vios.system.config.hostname
importvg	All	vios.lvm.manage.import
invscout	All	vios.system.firmware.scout
ioslevel	All	vios.system.level
ldapadd	All	vios.security.manage.ldap.add
ldapsearch	All	vios.security.manage.ldap.search
ldfware	All	vios.system.firmware.load
license	**-accept**	vios.system.license
	Others	vios.system.license.view
loadopt	All	vios.device.manage.optical.load
loginmsg	All	vios.security.user.login.msg
lsauth	All	vios.security.auth.list
lsdev	All	vios.device.manage.list
lsfailedlogin	All	vios.security.user.login.fail

Table 53. Authorizations corresponding to Virtual I/O Server commands (continued)

Command	Command options	Authorization
lsfware	All	vios.system.firmware.list
lsgcl	All	vios.security.log.list
lslparinfo	All	vios.system.lpar.list
lslv	All	vios.lvm.manage.list
lsmap	All	vios.device.manage.map.phyvirt
lsnetsvc	All	vios.network.service.list
lsnports	All	vios.device.manage.list
lspath	All	vios.device.manage.list
lspv	All	vios.device.manage.list
lsrep	All	vios.device.manage.repos.list
lsrole	All	vios.security.role.list
lssecattr	**-c**	vios.security.cmd.list
	-d	vios.security.device.list
	-f	vios.security.file.list
	-p	vios.security.proc.list
lssp	All	vios.device.manage.spool.list
lssvc	All	vios.system.config.agent.list
lssw	All	vios.system.software.list
lstcpip	All	vios.network.tcpip.list
lsuser	All	vios.security.user.list **Note:** Any user can run this command to view a minimal set of user attributes. However, only users with this authorization can view all the user attributes.
lsvg	All	vios.lvm.manage.list
lsvlog	All	vios.device.manage.vlog.list
lsvlrepo	All	vios.device.manage.vlrepo.list
lsvopt	All	vios.device.manage.optical.list
lu	**-create**	vios.device.manage.backing.create or vios.system.cluster.lu.create
	-map	vios.device.manage.backing.create or vios.system.cluster.lu.create or vios.system.cluster.lu.map
	-remove	vios.device.manage.backing.remove or vios.system.cluster.lu.remove
	-unmap	vios.device.manage.remove or vios.system.cluster.lu.unmap

Table 53. Authorizations corresponding to Virtual I/O Server commands (continued)		
Command	**Command options**	**Authorization**
migratepv	All	vios.device.manage.migrate
mirrorios	All	vios.lvm.manage.mirrorios.create
mkauth	All	vios.security.auth.create
mkbdsp	All	vios.device.manage.backing.create
mkkrb5clnt	All	vios.security.manage.kerberos.create
mkldap	All	vios.security.manage.ldap.create
mklv	All	vios.lvm.manage.create
mklvcopy	All	vios.lvm.manage.mirror.create
mkpath	All	vios.device.manage.path.create
mkrep	All	vios.device.manage.repos.create
mkrole	All	vios.security.role.create
mksp	All	vios.device.manage.spool.create
mktcpip	All	vios.network.tcpip.config
mkuser	All	vios.security.user.create
mkvdev	**-fbo**	vios.device.manage.create.virtualdisk
	-lnagg	vios.device.manage.create.lnagg
	-sea	vios.device.manage.create.sea
	-vdev	vios.device.manage.create.virtualdisk
	-vlan	vios.device.manage.create.vlan
mkvg	All	vios.lvm.manage.create
mkvlog	All	vios.device.manage.vlog.create
mkvopt	All	vios.device.manage.optical.create
motd	All	vios.security.user.msg
mount	All	vios.fs.mount
netstat	All	vios.network.tcpip.list
optimizenet	All	vios.network.config.tune
oem_platform_level	All	vios.system.level
oem_setup_env	All	vios.oemsetupenv
passwd	All	vios.security.passwd **Note:** A user can change the password without having this authorization. This authorization is required only if the user wants to change the password of other users.
pdump	All	vios.system.dump.platform
ping	All	vios.network.ping

Table 53. Authorizations corresponding to Virtual I/O Server commands (continued)

Command	Command options	Authorization
postprocesssvc	All	vios.system.config.agent
prepdev	All	vios.device.config.prepare
pv	**-add**, **-remove**, **-replace**	vios.device.manage.spool.change or vios.system.cluster.pool.modify
redefvg	All	vios.lvm.manage.reorg
reducevg	All	vios.lvm.manage.change
refreshvlan	All	vios.network.config.refvlan
remote_management	All	vios.system.manage.remote
replphyvol	All	vios.device.manage.replace
restore	All	vios.fs.backup
restorevgstruct	All	vios.lvm.manage.restore
rmauth	All	vios.security.auth.remove
rmbdsp	All	vios.device.manage.backing.remove
rmdev	All	vios.device.manage.remove
rmlv	All	vios.lvm.manage.remove
rmlvcopy	All	vios.lvm.manage.mirror.remove
rmpath	All	vios.device.manage.path.remove
rmrep	All	vios.device.manage.repos.remove
rmrole	All	vios.security.role.remove
rmsecattr	**-c**	vios.security.cmd.remove
	-d	vios.security.device.remove
	-f	vios.security.file.remove
rmsp	All	vios.device.manage.spool.remove
rmtcpip	All	vios.network.tcpip.remove
rmuser	All	vios.security.user.remove
rmvdev	All	vios.device.manage.remove
rmvlog	All	vios.device.manage.vlog.remove
rmvopt	All	vios.device.manage.optical.remove
rolelist	**-p**	vios.security.proc.role.list **Note:** You can run other options of this command without having any authorizations.
	-u	vios.security.role.list
savevgstruct	All	vios.lvm.manage.save
save_base	All	vios.device.manage.saveinfo
seastat	All	vios.network.stat.sea

Table 53. Authorizations corresponding to Virtual I/O Server commands (continued)

Command	Command options	Authorization
setkst	All	vios.security.kst.set
setsecattr	**-c**	vios.security.cmd.set
	-d	vios.security.device.set
	-f	vios.security.file.set
	-o	vios.security.domain.set
	-p	vios.security.proc.set
showmount	All	vios.fs.mount.show
shutdown	All	vios.system.boot.shutdown
snap	All	vios.system.trace.format
snapshot	All	vios.device.manage.backing.create
snmp_info	All	vios.security.manage.snmp.info
snmpv3_ssw	All	vios.security.manage.snmp.switch
snmp_trap	All	vios.security.manage.snmp.trap
startnetsvc	All	vios.network.service.start
startsvc	All	vios.system.config.agent.start
startsysdump	All	vios.system.dump
starttrace	All	vios.system.trace.start
stopnetsvc	All	vios.network.service.stop
stopsvc	All	vios.system.config.agent.stop
stoptrace	All	vios.system.trace.stop
svmon	All	vios.system.stat.memory
syncvg	All	vios.lvm.manage.sync
sysstat	All	vios.system.stat.list
rmsecattr	**-c**	vios.security.cmd.remove
	-d	vios.security.device.remove
	-f	vios.security.file.remove
tier	-create	vios.device.manage.spool.change or vios.system.cluster.pool.modify
	-remove	vios.device.manage.spool.change or vios.system.cluster.pool.modify
	-modify	vios.device.manage.spool.change or vios.system.cluster.pool.modify
topas	All	vios.system.config.topas
topasrec	All	vios.system.config.topasrec

Table 53. Authorizations corresponding to Virtual I/O Server commands (continued)

Command	Command options	Authorization
tracepriv	All	vios.security.priv.trace
traceroute	All	vios.network.route.trace
uname	All	vios.system.uname
unloadopt	All	vios.device.manage.optical.unload
unmirrorios	All	vios.lvm.manage.mirrorios.remove
unmount	All	vios.fs.unmount
updateios	All	vios.install
vasistat	All	vios.network.stat.vasi
vfcmap	All	vios.device.manage.map.virt
viosbr	**-view**	vios.system.backup.cfg.view
	Others	vios.system.backup.cfg **Note:** To run any other options of this command, this authorization is required.
viosecure	All	vios.security.manage.firewall
viostat	All	vios.system.stat.io
vmstat	All	vios.system.stat.memory
wkldagent	All	vios.system.manage.workload.agent
wkldmgr	All	vios.system.manage.workload.manager
wkldout	All	vios.system.manage.workload.process

Roles

The Virtual I/O Server retains its current roles and has the appropriate authorizations assigned to the roles. Additional roles that closely emulate the roles in the AIX operating system can be created. Additional roles that closely emulate the roles in the AIX operating system can be created. The roles emulate naming conventions and descriptions, but are only applicable to the Virtual I/O Server specific requirements. Users cannot view, use, or modify any of the default roles in the AIX operating system in the AIX operating system.

The following roles are the default roles in the AIX operating system in the AIX operating system. These roles are unavailable to the Virtual I/O Server users, and are not displayed.

- AccountAdmin
- BackupRestore
- DomainAdmin
- FSAdmin
- SecPolicy
- SysBoot
- SysConfig
- isso
- sa
- so

The following roles are the default roles in the Virtual I/O Server:

- Admin
- DEUser
- PAdmin
- RunDiagnostics
- SRUser
- SYSAdm
- ViewOnly

The **mkrole** command creates a role. The *newrole* parameter must be a unique role name. You cannot use the **ALL** or **default** keywords as the role name. Every role must have a unique role ID that is used for security decisions. If you do not specify the **id** attribute when you create a role, the **mkrole** command automatically assigns a unique ID to the role.

Naming convention: There is no standard naming convention for roles. However, existing names of roles cannot be used for creating roles.

Restriction:

The role parameter cannot contain spaces, tabs, or newline characters. To prevent inconsistencies, restrict role names to characters in the POSIX portable file name character set. You cannot use the keywords **ALL** or **default** as a role name. Do not use the following characters within a role-name string:

- : (colon)
- " (quotation mark)
- # (number sign)
- , (comma)
- = (equal sign)
- \ (backslash)
- / (forward slash)
- ? (question mark)
- ' (single quotation mark)
- ` (grave accent)

Privileges

A **Privilege** is an attribute of a process through which the process can bypass specific restrictions and limitations of the system. Privileges are associated with a process and are acquired by running a privileged command. Privileges are defined as bit-masks in the operating system kernel and enforce access control over privileged operations. For example, the privilege bit **PV_KER_TIME** might control the kernel operation to modify the system date and time. Nearly 80 privileges are included with the operating system and provide granular control over privileged operations. You can acquire the least privilege that is required to perform an operation through division of privileged operations in the kernel. This feature leads to enhanced security because a process hacker can get access to only 1 or 2 privileges in the system, and not to root user privileges.

Authorizations and roles are a user-level tool to configure user access to privileged operations. Privileges are the restriction mechanism that is used in the operating system kernel to determine whether a process has authorization to perform an action. Hence, if a user is in a role session that has an authorization to run a command, and that command is run, a set of privileges are assigned to the process. There is no direct mapping of authorizations and roles to privileges. Access to several commands can be provided through an authorization. Each of those commands can be granted a different set of privileges.

The following table lists the commands that are related to role-based access control (RBAC).

Table 54. RBAC commands and their descriptions

Command	Description
chauth	Modifies attributes of the authorization that is identified by the *newauth* parameter
chrole	Changes attributes of the role that is identified by the *role* parameter
lsauth	Displays attributes of user-defined and system-defined authorizations from the authorization database
lsrole	Displays the role attributes
lssecattr	Lists the security attributes of one or more commands, devices, or processes
mkauth	Creates new user-defined authorizations in the authorization database
mkrole	Creates new roles
rmauth	Removes the user-defined authorization that is identified by the *auth* parameter
rmrole	Removes the role that is identified by the *role* parameter from the roles database
rmsecattr	Removes the security attributes for a command, a device, or a file entry that is identified by the *Name* parameter from the appropriate database
rolelist	Provides role and authorization information to the caller about the roles that are assigned to them
setkst	Reads the security databases and loads the information from the databases into the kernel security tables
setsecattr	Sets the security attributes of the command, device, or process that are specified by the *Name* parameter
swrole	Creates a role session with the roles that are specified by the *Role* parameter
tracepriv	Records the privileges that a command attempts to use when the command is run

Managing users on the Virtual I/O Server

You can create, list, change, switch, and remove users by using Virtual I/O Server or the IBM Tivoli Identity Manager.

About this task

When the Virtual I/O Server is installed, the only user type that is active is the prime administrator (**padmin** having the default role **PAdmin**). The prime administrator can create additional user IDs with types of system administrator, service representative, development engineer, or other users with different roles.

Note: You cannot create the prime administrator (**padmin**) user ID. It is automatically created, enabled, and the role **PAdmin** is assigned as the default role after the Virtual I/O Server is installed.

The following table lists the user management tasks available on the Virtual I/O Server, as well as the commands you must run to accomplish each task.

Table 55. Tasks and associated commands for working with Virtual I/O Server users

Task	Command
Change passwords	**cfgassist**

Table 55. Tasks and associated commands for working with Virtual I/O Server users (continued)

Task	Command
Create a system administrator user ID	**mkuser**. This assigns **Admin** as the default role.
Create a service representative (SR) user ID	**mkuser** with the **-sr** flag. This assigns **SRUser** as the default role.
Create a development engineer (DE) user ID	**mkuser** with the **-de** flag. This assigns **DEUser** as the default role.
Create users with varied access rights	**mkuser** with the **-attr** flag by specifying **roles** and **default_roles** attributes. This assigns users with varied access rights, enabling them to access a varied set of commands.
Create an LDAP user	**mkuser** with the **-ldap** flag
List a user's attributes For example, determine whether a user is an LDAP user.	**lsuser**
Change a user's attributes	**chuser**
Switch to another user	**su**
Remove a user	**rmuser**

You can use the IBM Tivoli Identity Manager to automate the management of Virtual I/O Server users. Tivoli Identity Manager provides a Virtual I/O Server adapter that acts as an interface between the Virtual I/O Server and the Tivoli Identity Manager Server. The adapter acts as a trusted virtual administrator on the Virtual I/O Server, performing tasks like the following:

- Creating a user ID to authorize access to the Virtual I/O Server.
- Modifying an existing user ID to access the Virtual I/O Server.
- Removing access from a user ID. This deletes the user ID from the Virtual I/O Server.
- Suspending a user account by temporarily deactivating access to the Virtual I/O Server.
- Restoring a user account by reactivating access to the Virtual I/O Server.
- Changing a user account password on the Virtual I/O Server.
- Reconciling the user information of all current users on the Virtual I/O Server.
- Reconciling the user information of a particular user account on the Virtual I/O Server by performing a lookup.

For more information, see the IBM Tivoli Identity Manager product manuals.

Troubleshooting the Virtual I/O Server

Find information about diagnosing Virtual I/O Server problems and information about how to correct those problems.

Troubleshooting the Virtual I/O Server logical partition

Find information and procedures for troubleshooting and diagnosing the Virtual I/O Server logical partition.

Troubleshooting virtual SCSI problems

Find information and procedures for troubleshooting virtual Small Computer Serial Interface (SCSI) problems in the Virtual I/O Server.

About this task

For problem determination and maintenance, use the **diagmenu** command that is provided by the Virtual I/O Server.

If you are still having problems after you use the **diagmenu** command, contact your next level of support and ask for assistance.

Correcting a failed shared Ethernet adapter configuration

You can troubleshoot errors that occur when you configure a shared Ethernet adapter (SEA), such as errors that result in message 0514-040, by using the **lsdev**, **netstat**, and **entstat** commands.

Before you begin

When you configure an SEA, the configuration can fail with the following error:

```
Method error (/usr/lib/methods/cfgsea):
        0514-040 Error initializing a device into the kernel.
```

About this task

To correct the problem, complete the following steps:

Procedure

1. Verify that the physical and virtual adapters that are being used to create the shared Ethernet adapter are available by running the following command:

   ```
   lsdev -type adapter
   ```

2. Make sure that the interface of neither the physical nor any of the virtual adapters are configured. Run the following command:

   ```
   netstat -state
   ```

   Important: None of the interfaces of the adapters must be listed in the output. If any interface name (for example, *en0*) is listed in the output, detach it as follows:

   ```
   chdev -dev interface_name -attr state=detach
   ```

   You might want to perform this step from a console connection because it is possible that detaching this interface might end your network connection to the Virtual I/O Server.

3. Verify that the virtual adapters that are used for data are trunk adapters by running the following command:

```
entstat -all entX | grep Trunk
```

Note:

- The trunk adapter does not apply to the virtual adapter that is used as the control channel in an SEA Failover configuration.
- If any of the virtual adapters that are used for data are not trunk adapters, you need to enable them to access external networks from the HMC.

4. Verify that the physical device and the virtual adapters in the SEA agree on the checksum offload setting:

   a) Determine the checksum offload setting on physical device by running the following command:

   ```
   lsdev -dev device_name -attr chksum_offload
   ```

   where, *device_name* is the name of the physical device. For example, ent0.

   b) If `chksum_offload` is set to `yes`, enable checksum offload for all of the virtual adapters in the SEA by running the following command:

   ```
   chdev -dev device_name -attr chksum_offload=yes
   ```

   Where *device_name* is the name of a virtual adapter in the SEA. For example, ent2.

   c) If `chksum_offload` is set to no, disable checksum offload for all of the virtual adapters in the SEA by running the following command:

   ```
   chdev -dev device_name -attr chksum_offload=no
   ```

   where, *device_name* is the name of a virtual adapter in the SEA.

   d) If there is no output, the physical device does not support checksum offload and therefore does not have the attribute. To resolve the error, disable checksum offload for all of the virtual adapters in the SEA by running the following command:

   ```
   chdev -dev device_name -attr chksum_offload=no
   ```

   where, *device_name* is the name of a virtual adapter in the SEA.

5. If the real adapter is a logical host Ethernet adapter (LHEA) port, also known as a logical integrated virtual Ethernet adapter port, ensure that the Virtual I/O Server is configured as the promiscuous logical partition for the physical port of the logical integrated virtual Ethernet adapter from the HMC.

Debugging problems with Ethernet connectivity

You can determine Ethernet connectivity problems by examining Ethernet statistics that are produced by the **entstat** command. Then, you can debug the problems by using the **starttrace** and **stoptrace** commands.

About this task

To help debug problems with Ethernet connectivity, follow these steps:

Procedure

1. Verify that the source client logical partition can ping another client logical partition on the same system without going through the Virtual I/O Server.

   If this fails, the problem is likely in the client logical partition's virtual Ethernet setup. If the ping is successful, proceed to the next step.

2. Start a ping on the source logical partition to a destination machine so that the packets are sent through the Virtual I/O Server.

   This ping most likely fails. Proceed to the next step with the ping test running.

3. On the Virtual I/O Server, type the following command:

   ```
   entstat -all SEA_adapter
   ```

   where, *SEA_adapter* is the name of your Shared Ethernet Adapter.
4. Verify that the VLAN ID to which the logical partition belongs is associated with the correct virtual adapter in the VLAN IDs section of the output. Examine the `ETHERNET STATISTICS` for the virtual adapter for this VLAN and verify that the packet counts under the `Receive statistics` column are increasing.

   This verifies that the packets are being received by the Virtual I/O Server through the correct adapter. If the packets are not being received, the problem might be in the virtual adapter configuration. Verify the VLAN ID information for the adapters by using the Hardware Management Console (HMC).
5. Examine the `ETHERNET STATISTICS` for the physical adapter for this VLAN and verify that the packet counts under the `Transmit statistics` column are increasing.

   This step verifies that the packets are being sent out of the Virtual I/O Server.

   - If this count is increasing, then the packets are going out of the physical adapter. Continue to step “6” on page 254.
   - If this count is not increasing, then the packets are not going out of the physical adapter, and to further debug the problem, you must begin the system trace utility. Follow the instructions in step “9” on page 254 to collect a system trace, statistical information, and the configuration description. Contact service and support if you need to debug the problem further.
6. Verify that the target system outside (on physical side of Virtual I/O Server) is receiving packets and sending out replies.

   If this is not happening, either the wrong physical adapter is associated with the Shared Ethernet Adapter or the Ethernet switch might not be configured correctly.
7. Examine the `ETHERNET STATISTICS` for the physical adapter for this VLAN and verify that the packet counts under the Receive statistics column are increasing.

   This step verifies that the ping replies are being received by the Virtual I/O Server.

   If this count is not increasing, the switch might not be configured correctly.
8. Examine the `ETHERNET STATISTICS` for the virtual adapter for this VLAN and verify that the packet counts under the `Transmit statistics` column are increasing.

   This step verifies that the packet is being transmitted by the Virtual I/O Server through the correct virtual adapter.

   If this count is not increasing, start the system trace utility. Follow the instructions in step “9” on page 254 to collect a system trace, statistical information, and the configuration description. Work with service and support to debug the problem further.
9. Use the Virtual I/O Server trace utility to debug connectivity problems.

   Start a system trace by using the **starttrace** command specifying the trace hook ID. The trace hook ID for Shared Ethernet Adapter is 48F. Use the **stoptrace** command to stop the trace. Use the **cattracerpt** command to read the trace log, format the trace entries, and write a report to standard output.

Enabling noninteractive shells on Virtual I/O Server 1.3 or later

After you upgrade the Virtual I/O Server to 1.3 or later, you can enable noninteractive shells by using the **startnetsvc** command.

Before you begin

If you installed OpenSSH on a level of the Virtual I/O Server before 1.3, and then upgraded to 1.3, or later, noninteractive shells might not work because the SSH configuration file needs modification.

Procedure

To enable noninteractive shells in Virtual I/O Server 1.3, or later, run the following command from the SSH client:

```
ioscli startnetsvc ssh
```

Note: You can run the **startnetsvc** command when the SSH service is running. In this situation, the command appears to fail, but is successful.

Recovering when disks cannot be located

Learn how to recover from disks not displaying when trying to boot or install a client logical partition.

About this task

Occasionally, the disk that is needed to install the client logical partition cannot be located. In this situation, if the client is already installed, start the client logical partition. Ensure that you have the latest levels of the software and firmware. Then, ensure that the **Slot number** of the virtual Small Computer Serial Interface (SCSI) server adapter matches the **Remote partition virtual slot number** of the virtual SCSI client adapter.

When the Hardware Management Console (HMC) is at version 8.7.0, or later, if the storage has not been attached using the HMC interface, use the adapter view in the Virtual Storage of Manage Partition to verify the adapter mappings.

For more information about verifying adapter mappings when the HMC is at version 8.7.0, or later, see Managing virtual storage for a partition.

Troubleshooting AIX client logical partitions

Find information and procedures for troubleshooting AIX client logical partitions.

About this task

If your client partition is using virtual I/O resources, check the Service Focal Point and Virtual I/O Server first to ensure that the problem is not on the server.

On client partitions running the current level of AIX, when a hardware error is logged on the server and a corresponding error is logged on the client partition, the Virtual I/O Server provides a correlation error message in the error report.

Run the following command to gather an error report:

```
errpt -a
```

Running the **errpt** command returns results similar to the following:

```
LABEL:          VSCSI_ERR2
IDENTIFIER:     857033C6

Date/Time:       Tue Feb 15 09:18:11 2005
Sequence Number: 50
Machine Id:      00C25EEE4C00
Node Id:         vio_client53A
Class:           S
Type:            TEMP
Resource Name:   vscsi2

Description
Underlying transport error

Probable Causes
PROCESSOR

Failure Causes
PROCESSOR
```

```
        Recommended Actions
        PERFORM PROBLEM DETERMINATION PROCEDURES
Detail Data
Error Log Type
01
Reserve
00
Error Number
0006
RC
0000 0002
VSCSI Pointer
```

Compare the `LABEL`, `IDENTIFIER`, and `Error Number` values from your error report to the values in the following table to help identify the problem and determine a resolution.

Table 56. Labels, identifiers, error numbers, problem descriptions, and resolutions of common virtual Small Computer Serial Interface (SCSI) client logical partition problems

Label	Identifier	Error Number	Problem	Resolution
VSCSI_ERR2	857033C6	0006 RC 0000 0002	The virtual SCSI server adapter on the Virtual I/O Server logical partition is not open.	Make the server adapter on the Virtual I/O Server logical partition available for use.
		001C RC 0000 0000	The virtual SCSI server adapter on the Virtual I/O Server logical partition has been closed abruptly.	Determine why the server adapter in the Virtual I/O Server logical partition was closed.
VSCSI_ERR3	ED995F18	000D RC FFFF FFF0	The virtual SCSI server adapter on the Virtual I/O Server logical partition is being used by another client logical partition.	Terminate the client logical partition that is using the server adapter.
		000D RC FFFF FFF9	The virtual SCSI server adapter (partition number and slot number) specified in the client adapter definition does not exist.	On the HMC, correct the client adapter definition to associate it with a valid server adapter.

Performance data collection for analysis by the IBM Electronic Service Agent

You can use a number of Virtual I/O Server commands to collect various levels of performance data. This data can then be used by the IBM Electronic Service Agent support personnel to diagnose and solve performance problems.

The Virtual I/O Server Version 2.1.2.0 provides commands that you can use to capture performance data. You can then convert this data into a format and file for diagnostic use by the IBM Electronic Service Agent.

You can use the **cfgassist** command to manage the various types of data recording that the **topas** and **topasrec** commands provide. You can use the **wkldout** command to convert recording data from binary format to ASCII text format. You also can configure the performance management agent to gather data about performance of the Virtual I/O Server.

With the **topasrec** command, the Virtual I/O Server supports local, central electronics process (CEC), and cluster recording capabilities. These recordings can be either persistent or normal. Persistent recordings are recordings that run on the Virtual I/O Server and continue to run after the Virtual I/O Server reboots. Normal recordings are recordings that run for a specified time interval. The recording data files that are generated are stored in the `/home/ios/perf/topas` directory path.

Local recordings gather data about the Virtual I/O Server. CEC recordings gather data about any AIX logical partitions that are running on the same CEC as the Virtual I/O Server.CEC recordings gather data about any AIX logical partitions that are running on the same CEC as the Virtual I/O Server. The data that is collected consists of dedicated and shared logical partition data and includes a set of aggregated values that provide an overview of the partition set. Cluster recordings gather data from a list of hosts that are specified in a cluster configuration file.

The performance manager agent (named **perfmgr**) collects data about system performance and sends it to support through the Electronic Service Agent (ESA) for processing. When the agent is started, it runs a set of utilities to collect metrics to measure performance. After you configure the performance management agent, you can use the **startsvc**, **stopsvc**, **lssvc**, and **cfgsvc** commands to manage the agent. You can use the **postprocesssvc** command to generate an appropriately formatted file from a list of available individual performance data files. This file can then be understood by the Electronic Service Agent.

Related information

cfgassist command

cfgsvc command

lssvc command

postprocesssvc command

startsvc command

stopsvc command

topas command

topasrec command

wkldout command

Reference information for the Virtual I/O Server

Find reference information about the Virtual I/O Server commands, configuration attributes for Tivoli agents and clients, networking statistics and attributes, and Virtual I/O Server user types.

Virtual I/O Server command descriptions

You can view a description of each Virtual I/O Server command.

For more information about Virtual I/O Server commnds, see Virtual I/O Server commands.

Configuration attributes for IBM Tivoli agents and clients

Learn about required and optional configuration attributes and variables for the IBM Tivoli Monitoring agent, the IBM Tivoli Usage and Accounting Manager agent, the IBM Tivoli Storage Manager client, and the Tivoli Storage Productivity Center agents.

In the following tables, the term *attribute* refers to an option that you can add to a Virtual I/O Server command. The term *variable* refers to an option that you can specify in a configuration file for Tivoli Storage Manager or Tivoli Usage and Accounting Manager.

IBM Tivoli Monitoring

Table 57. Tivoli Monitoring configuration attributes

Attribute	Description
HOSTNAME	The host name or IP address of the Tivoli Enterprise Monitoring Server server to which the monitoring agent sends data.
MANAGING_SYSTEM	The host name or IP address of the Hardware Management Console (HMC) attached to the managed system on which the Virtual I/O Server with the monitoring agent is located. You can specify only one HMC per monitoring agent. If you do not specify the MANAGING_SYSTEM attribute, the Virtual I/O Server uses the Resource Monitoring and Control (RMC) connection to obtain the host name of IP address of the HMC.
RESTART_ON_REBOOT	Determines whether the monitoring agent restarts whenever the Virtual I/O Server restarts. TRUE indicates that the monitoring agent restarts whenever the Virtual I/O Server restarts. FALSE indicates that the monitoring agent does not restart whenever the Virtual I/O Server restarts.

IBM Tivoli Storage Manager

Table 58. Tivoli Storage Manager configuration attributes

Attribute	Description
SERVERNAME	The host name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.
SERVERIP	The IP address or domain name of the Tivoli Storage Manager server to which the Tivoli Storage Manager client is associated.
NODENAME	The name of the machine on which the Tivoli Storage Manager client is installed.

IBM Tivoli Usage and Accounting Manager

Table 59. Tivoli Usage and Accounting Manager configuration variables in the A_config.par file

Variable	Description	Possible values	Default value
AACCT_TRANS_IDS	Designates the AIXAIX advanced accounting record types that are included within the usage reports.	1, 4, 6, 7, 8, 10, 11, or 16	10

Table 59. Tivoli Usage and Accounting Manager configuration variables in the A_config.par file (continued)

Variable	Description	Possible values	Default value
AACCT_ONLY	Determines whether the Usage and Accounting Manager agent collects accounting data.	- Y: Indicates that the Usage and Accounting Manager agent collects accounting data. - N: Indicates that the Usage and Accounting Manager agent does not collect accounting data.	Y
ITUAM_SAMPLE	Determines whether the Usage and Accounting Manager agent collects data about the storage file system.	- Y: Indicates that the Usage and Accounting Manager agent collects data about the storage file system. - N: Indicates that the Usage and Accounting Manager agent does not collect data about the storage file system.	N

Table 60. Tivoli Usage and Accounting Manager configuration attributes

Attribute	Description
ACCT_DATA0	The size, in MB, of the first data file that holds daily accounting information.
ACCT_DATA1	The size, in MB, of the second data file that holds daily accounting information.
ISYSTEM	The time, in minutes, when the agent generates system interval records.
IPROCESS	The time, in minutes, when the system generates aggregate process records.

Tivoli Storage Productivity Center attributes

Table 61. Tivoli Storage Productivity Center configuration attributes

Attribute	Description	Required or optional
S	Host name or IP address of the Tivoli Storage Productivity Center Server that is associated with the Tivoli Storage Productivity Center agent.	Required
A	Host name or IP address of the Agent Manager.	Required

Table 61. Tivoli Storage Productivity Center configuration attributes (continued)

Attribute	Description	Required or optional
devAuth	Password for authentication to the Tivoli Storage Productivity Center device server.	Required
caPass	Password for authentication to the command agent.	Required
caPort	Number that identifies the port for the common agent. The default is 9510.	Optional
amRegPort	Number that identifies the registration port for the Agent Manager. The default is 9511.	Optional
amPubPort	Number that identifies the public port for the Agent Manager. The default is 9513.	Optional
dataPort	Number that identifies the port for the Tivoli Storage Productivity Center Data server. The default is 9549.	Optional
devPort	Number that identifies the port of the Tivoli Storage Productivity Center Device server. The default is 9550.	Optional
newCA	The default is true.	Optional
oldCA	The default is false.	Optional
daScan	Runs a scan for the TPC_data agent after installation. The default is true.	Optional
daScript	Runs the script for the TPC_data agent after installation. The default is true.	Optional
daIntsall	Installs the TPC_data agent. The default is true.	Optional
faInstall	Installs the TPC_fabric agent. The default is true.	Optional
U	Uninstalls the Tivoli Storage Productivity Center agents. Possible values include: • all • data • fabric	Optional

Related information

IBM Tivoli Application Dependency Discovery Manager Information Center

IBM Tivoli Identity Manager

IBM Tivoli Monitoring version 6.2.1 documentation
IBM Tivoli Monitoring Virtual I/O Server Premium Agent User's Guide
IBM Tivoli Storage Manager
IBM Tivoli Usage and Accounting Manager Information Center
IBM TotalStorage Productivity Center Information Center

GARP VLAN Registration Protocol statistics

Learn about Bridge Protocol Data Unit (BPDU), Generic Attribute Registration Protocol (GARP), and GARP VLAN Registration Protocol (GVRP) displayed by running the **`entstat -all`** command. You can also view examples.

BPDU refers to all protocol packets that are exchanged between the switch and the Shared Ethernet Adapter. The only bridge protocol currently available with the Shared Ethernet Adapter is GARP. GARP is a generic protocol that is used to exchange attribute information between two entities. The only type of GARP currently available on the Shared Ethernet Adapter is GVRP. With GVRP, the attributes that are exchanged are VLAN values.

BPDU statistics

The BPDU statistics include all BPDU packets that are sent or received.

Table 62. Descriptions of BPDU statistics

BPDU statistic	Description
Transmit	**Packets** Number of packets sent. **Failed packets** Number of packets that could not be sent (for example, packets that could not be sent because there was no memory to allocate the outgoing packet).
Receive	**Packets** Number of packets received. **Unprocessed Packets** Packets that could not be processed because the protocol was not running at the time. **Non-contiguous Packets** Packets that were received in several packet fragments. **Packets with unknown PID** Packets that had a protocol ID (PID) different than GARP. A high number is typical because the switch might be exchanging other BPDU protocol packets that the Shared Ethernet Adapter does not support. **Packets with Wrong Length** Packets whose specified length (in the Ethernet header) does not match the length of the Ethernet packet received.

GARP statistics

The GARP statistics include those BPDU packets sent or received that are of type GARP.

Table 63. Descriptions of GARP statistics

GARP statistic	Description
Transmit	**Packets** Number of packets sent. **Failed packets** Number of packets that could not be sent (for example, packets that could not be sent because there was no memory to allocate the outgoing packet). **Leave All Events** Packets that are sent with event type *Leave All*. **Join Empty Events** Packets that are sent with event type *Join Empty* **Join In Events** Packets that are sent with event type *Join In* **Leave Empty Events** Packets that are sent with event type *Leave Empty* **Leave In Events** Packets that are sent with event type *Leave In* **Empty Events** Packets that are sent with event type *Empty*

Table 63. Descriptions of GARP statistics (continued)

GARP statistic	Description
Receive	**Packets** Number of packets received **Unprocessed Packets** Packets that could not be processed because the protocol was not running at the time. **Packets with Unknown Attr Type:** Packets with an unsupported attribute type. A high number is typical because the switch might be exchanging other GARP protocol packets that the Shared Ethernet Adapter does not support. For example, GARP Multicast Registration Protocol (GMRP). **Leave All Events** Packets that are received with event type *Leave All* **Join Empty Events** Packets that are received with event type *Join Empty* **Join In Events** Packets that are received with event type *Join In* **Leave Empty Events** Packets that are received with event type *Leave Empty* **Leave In Events** Packets that are received with event type *Leave In* **Empty Events** Packets that are received with event type *Empty*

GVRP statistics

The GVRP statistics include those GARP packets sent or received that are exchanging VLAN information by using GVRP.

Table 64. Descriptions of GVRP statistics

GVRP statistic	**Description**
Transmit	**Packets** Number of packets sent **Failed packets** Number of packets that could not be sent (for example, packets that could not be sent because there was no memory to allocate the outgoing packet). **Leave All Events** Packets that are sent with event type *Leave All*. **Join Empty Events** Packets that are sent with event type *Join Empty* **Join In Events** Packets that are sent with event type *Join In* **Leave Empty Events** Packets that are sent with event type *Leave Empty* **Leave In Events** Packets that are sent with event type *Leave In* **Empty Events** Packets that are sent with event type *Empty*

Table 64. Descriptions of GVRP statistics (continued)

GVRP statistic	Description
Receive	**Packets** Number of packets received. **Unprocessed Packets** Packets that could not be processed because the protocol was not running at the time. **Packets with Invalid Length** Packets that contain one or more attributes whose length does not correspond to its event type. **Packets with Invalid Event** Packets that contain one or more attributes whose event type is invalid. **Packets with Invalid Value** Packets that contain one or more attributes whose value is invalid (for example, an invalid VLAN ID). **Total Invalid Attributes** Sum of all of the attributes that had an invalid parameter. **Total Valid Attributes** Sum of all of the attributes that had no invalid parameters. **Leave All Events** Packets that are sent with event type *Leave All*. **Join Empty Events** Packets that are sent with event type *Join Empty* **Join In Events** Packets that are sent with event type *Join In* **Leave Empty Events** Packets that are sent with event type *Leave Empty* **Leave In Events** Packets that are sent with event type *Leave In* **Empty Events** Packets that are sent with event type *Empty*

Example statistics

Running the **`entstat -all`** command returns results similar to the following:

```
-------------------------------------------------------------
Statistics for adapters in the Shared Ethernet Adapter ent3
-------------------------------------------------------------
Number of adapters: 2
SEA Flags: 00000009
    < THREAD >
    < GVRP >
VLAN IDs :
    ent2: 1
Real Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets transmitted: 0
    Packets dropped: 0
Virtual Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets transmitted: 0
    Packets dropped: 0
Other Statistics:
    Output packets generated: 0
    Output packets dropped: 0
    Device output failures: 0
    Memory allocation failures: 0
    ICMP error packets sent: 0
    Non IP packets larger than MTU: 0
    Thread queue overflow packets: 0

--------------------------------------------------------------
Bridge Protocol Data Units (BPDU) Statistics:

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 2                                    Packets: 1370
Failed packets: 0                             Unprocessed Packets: 0
                                              Non-contiguous Packets: 0
                                              Packets w/ Unknown PID: 1370
                                              Packets w/ Wrong Length: 0

---------------------------------------------------------------
General Attribute Registration Protocol (GARP) Statistics:

Transmit Statistic:                           Receive Statistics:
-------------------                           -------------------
Packets: 2                                    Packets: 0
Failed packets: 0                             Unprocessed Packets: 0
                                              Packets w/ Unknow Attr. Type: 0

Leave All Events: 0                           Leave All Events: 0
Join Empty Events: 0                          Join Empty Events: 0
Join In Events: 2                             Join In Events: 0
Leave Empty Events: 0                         Leave Empty Events: 0
Leave In Events: 0                            Leave In Events: 0
Empty Events: 0                               Empty Events: 0

---------------------------------------------------------------
GARP VLAN Registration Protocol (GVRP) Statistics:

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 2                                    Packets: 0
Failed packets: 0                             Unprocessed Packets: 0
                                              Attributes w/ Invalid Length: 0
                                              Attributes w/ Invalid Event: 0
                                              Attributes w/ Invalid Value: 0
                                              Total Invalid Attributes: 0
                                              Total Valid Attributes: 0

Leave All Events: 0                           Leave All Events: 0
Join Empty Events: 0                          Join Empty Events: 0
Join In Events: 2                             Join In Events: 0
Leave Empty Events: 0                         Leave Empty Events: 0
Leave In Events: 0                            Leave In Events: 0
Empty Events: 0                               Empty Events: 0
```

Network attributes

Find instructions for managing network attributes.

You can use several of the Virtual I/O Server (VIOS) commands, including **chdev**, **mkvdev**, and **cfglnagg**, to change device or network attributes. This section defines attributes that can be modified.

Ethernet attributes

You can modify the following Ethernet attributes.

Attribute	Description
Maximum Transmission Unit (*mtu*)	Specifies maximum transmission unit (MTU). This value can be any number from 60 through 65535, but it is media-dependent.
Interface State (*state*)	**detach** Removes an interface from the network interface list. If the last interface is detached, the network interface driver code is unloaded. To change the interface route of an attached interface, that interface must be detached and added again with the **chdev -dev** *Interface* **-attr** *state=detach* command. **down** Marks an interface as inactive, which keeps the system from trying to transmit messages through that interface. However, routes that use the interface are not automatically disabled. (**chdev -dev** *Interface* **-attr** *state=down*) **up** Marks an interface as active. This parameter is used automatically when setting the first address for an interface. It can also be used to enable an interface after the **chdev -dev** *Interface* **-attr** *state=up* command.
Network Mask (*netmask*)	Specifies how much of the address to reserve for subdividing networks into subnetworks. The *mask* includes both the network part of the local address and the subnet part, which is taken from the host field of the address. The mask can be specified as a single hexadecimal number beginning with 0x, in standard internet dotted decimal notation. In the 32-bit address, the mask contains bits with a value of 1 for the bit positions that are reserved for the network and subnet parts, and a bit with the value of 0 for the bit positions that specify the host. The mask contains the standard network portion, and the subnet segment is contiguous with the network segment.

Shared Ethernet Adapter attributes

You can modify the following Shared Ethernet Adapter attributes.

Attribute	Description
PVID (*pvid*)	Port VLAN ID (PVID). Specifies the PVID to use for the Shared Ethernet Adapter. PVID specifies the VLAN ID that is used for the non-VLAN tagged packets. PVID must match the PVID of the adapter that is specified in the *pvid_adapter* attribute. The PVID of trunk adapters other than the default virtual adapter (*pvid_adapter*), cannot be used by any client LPARs. This is because packets that have the PVID of other trunk adapters, instead of the PVID of the default virtual adapter, have their VLAN tag removed and sent out as untagged packets to comply with the IEEE VLAN specification.
PVID adapter (*pvid_adapter*)	Specifies the default virtual adapter to use for non-VLAN tagged packets. PVID of the *pvid_adapter* attribute must be specified as the value for the pvid attribute.
Physical adapter (*real_adapter*)	Specifies the physical adapter that is associated with the Shared Ethernet Adapter.

Attribute	Description
Thread (*thread*)	Activates or deactivates threading on the Shared Ethernet Adapter. Activating this option adds approximately 16 - 20% more machine cycles per transaction for MTU 1500 streaming, and approximately 31 - 38% more machine cycles per transaction for MTU 9000. The threading option adds more machine cycles per transaction at lesser workloads due to the threads being started for each packet. At higher workload rates, such as full duplex or the request/response workloads, the threads can run longer without waiting and being redispatched. Threaded mode must be used when virtual Small Computer Serial Interface (SCSI) will be run on the same Virtual I/O Server logical partition as Shared Ethernet Adapter. Threaded mode helps ensure that virtual SCSI and the Shared Ethernet Adapter can share the processor resource appropriately. However, threading adds more instruction path length, which uses additional processor cycles. If the Virtual I/O Server logical partition will be dedicated to running shared Ethernet devices (and associated virtual Ethernet devices) only, the adapters must be configured with threading disabled. You can enable or disable threading using the **-attr thread** option of the **`mkvdev`** command. To enable threading, use the `-attr thread=1` option. To disable threading, use the `-attr thread=0` option. For example, the following command disables threading for Shared Ethernet Adapter `ent1`: `mkvdev -sea ent1 -vadapter ent5 -default ent5 -defaultid 1 -attr thread=0`
Virtual adapters (*virt_adapter*)	Lists the virtual Ethernet adapters that are associated with the Shared Ethernet Adapter.
TCP segmentation offload (*largesend*)	Enables TCP largesend capability (also known as segmentation offload) from logical partitions to the physical adapter. The physical adapter must be enabled for TCP largesend for the segmentation offload from the logical partition to the Shared Ethernet Adapter to work. Also, the logical partition must be capable of performing a largesend operation. On AIX, largesend can be enabled on a logical partition by using the **`ifconfig`** command.On AIX, largesend can be enabled on a logical partition by using the **`ifconfig`** command. You can enable or disable TCP largesend by using the `-a largesend` option of the **chdev** command. To enable it, use the '-a largesend=1' option. To disable it, use the '-a largesend=0' option. For example, the following command enables *largesend* for Shared Ethernet Adapter ent1: `chdev -l ent1 -a largesend=1` By default the setting is disabled (largesend=0). **Note:** Largesend is enabled by default (largesend=1) on VIOS 2.2.3.0 and higher. For VIOS 2.2.3.0 and higher, network interface that is configured over Shared Ethernet Adapter device supports largesend operation.

Attribute	Description
TCP large receive offload (*large_receive*)	Enables the TCP large receive offload capability on the real adapter. When it is set and if the real adapter supports it, packets received by the real adapter is aggregated before they are passed to the upper layer, resulting in better performance. This parameter must be enabled only if all the partitions that are connected to the shared Ethernet adapter can handle packets larger than their MTU. This is not the same for Linux partitions. If all the logical partitions that are connected to the shared Ethernet adapter are AIX systems, this parameter can be enabled.
Jumbo frames (*jumbo_frames*)	Allows the interface that is configured over the Shared Ethernet Adapter to increase its MTU to 9000 bytes (the default is 1500). If the underlying physical adapter does not support jumbo frames and the *jumbo_frames* attribute is set to `yes`, then configuration fails. The underlying physical adapter must support jumbo frames. The Shared Ethernet Adapter automatically enables jumbo frames on its underlying physical adapter if *jumbo_frames* is set to `yes`. You cannot change the value of *jumbo_frames* at run time.
GARP VLAN Registration Protocol (GVRP) (*gvrp*)	Enables and disables GVRP on a Shared Ethernet Adapter.
Quality of service (*qos_mode*)	Allows the shared Ethernet adapter to prioritize the traffic based on the IEEE 802.1Q (VLAN) Priority code point. When it is disabled, VLAN traffic is not inspected for priority and all frames are treated equally. In *strict* mode, the high priority traffic is sent preferentially over less priority traffic. This mode provides better performance and more bandwidth to more important traffic. This can result in substantial delays for less priority traffic. In *loose* mode, a cap is placed on each priority level so that after a number of bytes is sent for each priority level, the following level is serviced. This method ensures that all packets are eventually sent. The high priority traffic is given less bandwidth with this mode than with strict mode. The caps in *loose* mode are such that more bytes are sent for the high priority traffic, so it gets more bandwidth than less priority traffic.
Number of threads (*nthreads*)	Specifies the number of threads in threaded mode, where the value of the **`thread`** parameter is 1. This value applies only when the thread mode is enabled. The **`nthreads`** attribute can be set to any value in the range 1 - 128 and has a default value of 7.
Queue size (*queue_size*)	Specifies the queue size for the Shared Ethernet Adapter threads in threaded mode where the value of the **`thread`** parameter is 1. This attribute indicates the number of packets that can be accommodated in each thread queue. This value applies only when the thread mode is enabled. When you change this value, the change does not take effect until the system restarts.
Hash algorithms (*hash_algo*)	Specifies the hash algorithm that is used to assign connections to Shared Ethernet Adapter threads in threaded mode, where the value of the **`thread`** parameter is 1. When the **`hash_algo`** parameter is set to 0, an addition operation of the source and destination Media Access Control (MAC) addresses, IP addresses, and port numbers is done. When the **`hash_algo`** parameter is set to 1, a **`murmur3 hash`** function is done instead of an addition operation. The **`murmur3 hash`** function is slower, but it achieves better distribution. This value applies only when the thread mode is enabled.

Attribute	Description
Virtual server network (VSN) (*lldpsvc*)	Activates the VSN capability on the Shared Ethernet Adapter when you set the attribute to `yes`. The VSN capability can be enabled on the Hardware Management Console (HMC) Version 7 Release 7.7.0, or later. The default value of the **`lldpsvc`** attribute is no. This attribute must be set to no before you remove the Shared Ethernet Adapter. For example, the following command enables the VSN capability for the Shared Ethernet Adapter *ent1*: `chdev -dev ent1 -a lldpsvc=yes`
Accounting (*accounting*)	When enabled, the Shared Ethernet Adapter (SEA) keeps a count of the number of bytes and packets that are bridged to and from each client LPAR. Use the **`seastat`** command to see those statistics.
Detect flip flops (*ff_detect*)	When enabled, the system can detect flip flops. By default, this setting is disabled. *Flip flop* indicates a situation in which two SEAs are constantly switching between failover and failback events.
Flip flops action (*ff_action*)	When enabled, you can specify what action the system must take when a flip flop state is detected. This attribute is not supported when the **ff_detect** attribute is disabled. The **ff_action** attribute can have the following values: **standby** - Specifies that the SEA must be placed in the standby mode. You can use this mode to manually fix the SEA related system issues. **recover** - Specifies that the SEA must recover by itself. Only an SEA that is of higher priority detects the flip flop state and takes subsequent actions. The SEA changes into a flip flop state when the following conditions are met: • If the **ff_detect** attribute is enabled. • During the time interval of 20 + fb_delay seconds, if the SEA that is of higher priority becomes the primary SEA three or more times. **Note:** The traffic will not be bridged when the SEA is in a flip flop state. When a flip flop state is detected, and if the value of the **ff_action** attribute is set to **standby**, the SEA goes into a standby mode and you can manually fix the SEA related system issues. If the **ff_action** attribute is set to **recover**, the system will try to recover by itself. During the time interval that is set in the **health_time** attribute, the SEA remains in a flip flop state while monitoring the link status and keep alive packets. If the link remains up and if the keep alive packets are received regularly, the **recover** action starts after the time interval that is set in the **health_time** attribute has lapsed. If the SEA receives keep alive packets from an SEA that is of lower priority, it becomes a primary SEA. If the SEA receives keep alive packets from an SEA that is of higher priority, it becomes a backup SEA. **Note:** If the value of the **health_time** attribute is 0, the **recover** action will be attempted immediately without monitoring the link status and keep alive packets.

Attribute	Description
Platform Large Send (*plso_bridge*)	One of the requirements for a Linux client is that the Maximum Segment Size (MSS) value must be known to receive large send packets. When the **plso_bridge** attribute is enabled, and when large send packets are received by the SEA, the Shared Ethernet Adapter can convey the MSS values to the Linux client through the receive descriptor in the hypervisor. By default, the **plso_bridge** attribute is enabled.

Shared Ethernet Adapter failover attributes

You can modify the following Shared Ethernet Adapter failover attributes.

Attribute	Description
High availability mode (*ha_mode*)	Determines whether the devices participate in a failover setup. The default is `disabled`. Typically, a Shared Ethernet Adapter in a failover setup is operating in `auto` mode, and the primary adapter is decided based on which adapter has the highest priority (lowest numerical value). A shared Ethernet device can be forced into the standby mode, where it behaves as the backup device while it can detect the presence of a functional primary. The following are the possible values for the **High availability mode** attribute: **Disabled** This value is the default value. It indicates that the Shared Ethernet Adapter does not participate in Shared Ethernet Adapter failover configuration. You must use this value only if you do not want to use Shared Ethernet Adapter failover configuration on the system. **Restriction:** If the Shared Ethernet Adapter is configured previously in the Shared Ethernet Adapter failover configuration, do not use this value. **Auto** This value indicates that the Shared Ethernet Adapter is in traditional failover configuration. In this configuration, one Shared Ethernet Adapter is the primary adapter and the other Shared Ethernet Adapter is the backup adapter. Depending on the priority value of the trunk adapters, a Shared Ethernet Adapter is configured as the primary or the backup adapter. **Standby** A shared Ethernet device can be forced into the *Standby* mode. A device that is in this mode functions as the backup device for the duration in which it can detect a functional primary adapter. **Sharing** This value indicates that the Shared Ethernet Adapter is participating in load sharing. For the Shared Ethernet Adapter to participate in load sharing, the load sharing criteria must be met. Also, the **High availability mode** attribute must be set to the *Sharing* mode on both **Shared Ethernet Adapters**.
Control Channel (*ctl_chan*)	Sets the virtual Ethernet device that is required for a Shared Ethernet Adapter in a failover setup so that it can communicate with the other adapter. There is no default value for this attribute, and it is required when the *ha_mode* is not set to `disabled`. **Note:** The *Control Channel* attribute is an optional attribute with the Power Hypervisor Version 780, or later and with the VIOS Version 2.2.3.0, or later.

Attribute	Description
Internet address to ping (*netaddr*)	Optional attribute that can be specified for a Shared Ethernet Adapter that has been configured in a failover setup. When this attribute is specified, a shared Ethernet device will periodically ping the IP address to verify connectivity (in addition to checking for link status of the physical devices). If it detects a loss of connectivity to the specified ping host, it will initiate a failover to the backup Shared Ethernet Adapter. This attribute is not supported when you use a Shared Ethernet Adapter with a Host Ethernet Adapter (or Integrated Virtual Ethernet).
Adapter reset (*adapter_reset*)	When enabled, the shared Ethernet adapter disables and reenables its physical adapter whenever it becomes inactive. It might help the external switch to direct the traffic to the new server. By default the setting is disabled.
Enable Reverse ARP transmit (*send_RARP*)	When enabled, the shared Ethernet adapter sends a reverse ARP after the Shared Ethernet Adapter failover. The reverse ARP is sent by a new primary Shared Ethernet Adapter to notify the switches of routing change. By default, the setting is enabled.
Health Time (*health_time*)	Sets the time that is required to elapse before a system is considered "healthy" after a system failover. After a Shared Ethernet Adapter moves to an "unhealthy" state, the *Health Time* attribute specifies an integer that indicates the number of seconds for which the system must maintain a "healthy" state before it is allowed to return into the Shared Ethernet Adapter protocol. You can use the following command to display the default values for this attribute: `lsattr -D -c adapter -s pseudo -t sea -a health_time`
Link Time (*link_time*)	**Note:** Currently, the link status check is effectively disabled in levels that contain this fix due to APAR IV97991.
Failback delay (*fb_delay*)	Sets the time that is required to elapse before a higher priority Shared Ethernet Adapter begins the failback process to take over as the primary SEA after a failover event. The *Failback delay* attribute is a dynamic attribute that can be changed at run time. The new value governs the time delay in subsequent failover/failback events. You can use the following command to display the default values for this attribute: `lsattr -D -c adapter -s pseudo -t sea -a fb_delay`
No automatic failback (*noauto_failback*)	When enabled, the higher priority Shared Ethernet Adapter does not attempt to automatically take over the system after a failover event. Instead, it remains as the backup Shared Ethernet Adapter. When the *No automatic failback* attribute is disabled, the higher priority SEA begins the failback process to take over as the primary SEA. This attribute can be changed during run time. The change affects the behavior of the Shared Ethernet Adapter for subsequent failover/failback events. By default, this attribute is disabled.

INET attributes

You can modify the following INET attributes.

Attribute	Description
Host Name (*hostname*)	Specify the host name that you want to assign to the current machine. When specifying the host name, use ASCII characters, preferably alphanumeric only. Do not use a period in the host name. Avoid using hexadecimal or decimal values as the first character (for example `3Comm`, where 3C might be interpreted as a hexadecimal character). For compatibility with earlier hosts, use an unqualified host name of fewer than 32 characters. If the host uses a domain name server for name resolution, the host name must contain the full domain name. In the hierarchical domain naming system, names consist of a sequence of subnames that are not case-sensitive and that are separated by periods with no embedded blanks. The DOMAIN protocol specifies that a local domain name must be fewer than 64 characters, and that a host name must be fewer than 32 characters in length. The host name is given first. Optionally, the full domain name can be specified; the host name is followed by a period, a series of local domain names separated by periods, and finally by the root domain. A fully specified domain name for a host, including periods, must be fewer than 255 characters in length and in the following form: `host.subdomain.subdomain.rootdomain` In a hierarchical network, certain hosts are designated as name servers that resolve names into internet addresses for other hosts. This arrangement has two advantages over the flat name space: resources of each host on the network are not consumed in resolving names, and the person who manages the system does not need to maintain name-resolution files on each machine on the network. The set of names that are managed by a single name server is known as its *zone of authority*.
Gateway (*gateway*)	Identifies the gateway to which packets are addressed. The *Gateway* parameter can be specified either by symbolic name or numeric address.
Route (*route*)	Specifies the route. The format of the *Route* attribute is: *route=destination*, *gateway*, [*metric*]. **destination** Identifies the host or network to which you are directing the route. The *Destination* parameter can be specified either by symbolic name or numeric address. **gateway** Identifies the gateway to which packets are addressed. The *Gateway* parameter can be specified either by symbolic name or numeric address. **metric** Sets the routing metric. The default is 0 (zero). The routing metric is used by the routing protocol (the *routed* daemon). Higher metrics have the effect of making a route less favorable. Metrics are counted as additional hops to the destination network or host.

Adapter attributes

You can modify the following adapter attributes. The attribute behavior can vary, based on the adapter and driver you have.

Attribute	Adapters/Drivers	Description
Media Speed (*media_speed*)	• 2-Port 10/100/1000 Base-TX PCI-X Adapter • 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver	The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 10 Mbps half-duplex, 10 Mbps full-duplex, 100 Mbps half-duplex, 100 Mbps full-duplex and autonegotiation, with a default of autonegotiation. Select auto-negotiate when the adapter should use autonegotiation across the network to determine the speed. When the network will not support autonegotiation, select the specific speed. 1000 MBps half and full duplex are not valid values. According to the IEEE 802.3z specification, gigabit speeds of any duplexity must be autonegotiated for copper (TX)-based adapters. If these speeds are required, select auto-negotiate.
Media Speed (*media_speed*)	• 2-Port Gigabit Ethernet-SX PCI-X Adapter • Gigabit Ethernet-SX PCI-X Adapter Device Driver	The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 1000 Mbps full-duplex and autonegotiation. The default is autonegotiation. Select auto-negotiate when the adapter should use autonegotiation across the network to determine the duplexity. When the network does not support autonegotiation, select 1000 Mbps full-duplex.
Media Speed (*media_speed*)	• 10/100 Mbps Ethernet PCI Adapter Device Driver	The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 10 Mbps half-duplex, 10 Mbps full-duplex, 100 Mbps half-duplex, 100 Mbps full-duplex and autonegotiation, with a default of autonegotiation. When the adapter should use autonegotiation across the network to determine the speed, select autonegotiate. When the network will not support autonegotiation, select the specific speed. If autonegotiation is selected, the remote link device must also be set to autonegotiate to ensure the link works correctly.

Attribute	Adapters/Drivers	Description
Media Speed (*media_speed*)	- 10/100/1000 Base-T Ethernet PCI adapter - Gigabit Ethernet-SX PCI Adapter Device Driver	The media speed attribute indicates the speed at which the adapter attempts to operate. The available speeds are 10 Mbps half-duplex, 10 Mbps full-duplex, 100 Mbps half-duplex, 100 Mbps full-duplex and autonegotiation, with a default of autonegotiation. Select autonegotiate when the adapter should use autonegotiation across the network to determine the speed. When the network will not support autonegotiation, select the specific speed. For the adapter to run at 1000 Mbit/s, the autonegotiation setting must be selected. **Note:** For the Gigabit Ethernet-SX PCI Adapter, the only selection available is autonegotiation.
Enable Alternate Ethernet Address (*use_alt_addr*)		Setting this attribute to yes indicates that the address of the adapter, as it appears on the network, is the one specified by the Alternate Ethernet Address attribute. If you specify the no value, the unique adapter address written in a ROM on the adapter card is used. The default value is no.
Alternate Ethernet Address (*alt_addr*)		Allows the adapter unique address, as it appears on the LAN network to be changed. The value entered must be an Ethernet address of 12 hexadecimal digits and must not be the same as the address of any other Ethernet adapter. There is no default value. This field has no effect unless the Enable Alternate Ethernet Address attribute is set to yes value, in which case this field must be filled in. A typical Ethernet address is 0x02608C000001. All 12 hexadecimal digits, including leading zeros, must be entered.
Enable Link Polling (*poll_link*)	- 10/100Mbps Ethernet PCI Adapter Device Driver	Select no to cause the device driver to poll the adapter to determine the status of the link at a specified time interval. The time interval value is specified in the **Poll Link Time Interval** field. If you select no, the device driver will not poll the adapter for its link status. The default value is no.

Attribute	Adapters/Drivers	Description
Poll Link Time Interval (*poll_link_time*)	• 10/100Mbps Ethernet PCI Adapter Device Driver	The amount of time, in milliseconds, between polls to the adapter for its link status that the device driver is allowed. This value is required when the **Enable Link Polling** option is set to yes. A value between 100 through 1000 can be specified. The incremental value is 10. The default value is 500.
Flow Control (*flow_ctrl*)	• 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • Gigabit Ethernet-SX PCI Adapter Device Driver	This attribute specifies whether the adapter should enable transmit and receive flow control. The default value is no.
Transmit Jumbo Frames (*jumbo_frames*)	• 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • Gigabit Ethernet-SX PCI Adapter Device Driver	Setting this attribute to yes indicates that frames up to 9018 bytes might be transmitted on this adapter. If you specify no, the maximum size of frames that are transmitted is 1518 bytes. Frames up to 9018 bytes can always be received on this adapter.

Attribute	Adapters/Drivers	Description
Checksum Offload (*chksum_offload*)	• 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • Gigabit Ethernet-SX PCI Adapter Device Driver • Virtual Ethernet adapters	Setting this attribute to `yes` indicates that the adapter calculates the checksum for transmitted and received TCP frames. If you specify no, the checksum will be calculated by the appropriate software. When a virtual Ethernet adapter has checksum offload enabled, the adapter advertises it to the hypervisor. The hypervisor tracks which virtual Ethernet adapters have checksum offload enabled and manages inter-partition communication accordingly. When network packets are routed through the Shared Ethernet Adapter, there is a potential for link errors. In this environment, the packets must traverse the physical link with a checksum. Communication works in the following way: • When a packet is received from the physical link, the physical adapter verifies the checksum. If the packet's destination is a virtual Ethernet adapter with checksum offload enabled, the receiver does not have to perform checksum verification. A receiver that does not have checksum offload enabled will accept the packet after checksum verification. • When a packet originates from a virtual Ethernet adapter with checksum offload enabled, it travels to the physical adapter without a checksum. The physical adapter generates a checksum before sending the packet out. Packets originating from a virtual Ethernet adapter with checksum offload disabled generate the checksum at the source. To enable checksum offload for a Shared Ethernet Adapter, all constituent devices must have it enabled as well. The shared Ethernet device fails if the underlying devices do not have the same checksum offload settings.

Attribute	Adapters/Drivers	Description
Enable Hardware Transmit TCP Resegmentation (*large_send*)	• 10/100/1000 Base-T Ethernet PCI-X Adapter Device Driver • Gigabit Ethernet-SX PCI-X Adapter Device Driver • 2-Port 10/100/1000 Base-TX PCI-X Adapter • 2-Port Gigabit Ethernet-SX PCI-X Adapter • Gigabit Ethernet-SX PCI Adapter Device Driver	This attribute specifies whether the adapter is to perform transmit TCP resegmentation for TCP segments. The default value is no.

Link Aggregation (Etherchannel) device attributes

You can modify the following Link Aggregation, or Etherchannel attributes.

Attribute	Description
Link Aggregation adapters (*adapter_names*)	The adapters that currently make up the Link Aggregation device. If you want to modify these adapters, modify this attribute and select all the adapters that must belong to the Link Aggregation device. When you use this attribute to select all of the adapters that must belong to the Link Aggregation device, its interface must not have an IP address configured.
Mode (*mode*)	The type of channel that is configured. In standard mode, the channel sends the packets to the adapter based on an algorithm (the value that is used for this calculation is determined by the Hash Mode attribute). In round_robin mode, the channel gives one packet to each adapter before repeating the loop. The default mode is standard. Using the 802.3ad mode, the Link Aggregation Control Protocol (LACP) negotiates the adapters in the Link Aggregation device with an LACP-enabled switch. If the Hash Mode attribute is set to anything other than the default, this attribute must be set to standard or 802.3ad. Otherwise, the configuration of the Link Aggregation device fails.

Attribute	Description
Hash Mode (*hash_mode*)	If operating under standard or IEEE 802.3ad mode, the hash mode attribute determines how the outgoing adapter for each packet is chosen. Following are the different modes: • `default`: uses the destination IP address to determine the outgoing adapter. • `src_port`: uses the source TCP or UDP port for that connection. • `dst_port`: uses the destination TCP or UDP port for that connection. • `src_dst_port`: uses both the source and destination TCP or UDP ports for that connection to determine the outgoing adapter. You cannot use round-robin mode with any hash mode value other than default. The Link Aggregation device configuration fails if you attempt this combination. If the packet is not TCP or UDP, it uses the default hashing mode (destination IP address). Using TCP or UDP ports for hashing can make better use of the adapters in the Link Aggregation device because connections to the same destination IP address can be sent over different adapters (while still retaining the order of the packets), thus increasing the bandwidth of the Link Aggregation device.
Internet Address to Ping (*netaddr*)	This field is optional. The IP address that the Link Aggregation device should ping to verify that the network is up. This is only valid when there is a backup adapter and when there are one or more adapters in the Link Aggregation device. An address of zero (or all zeros) is ignored and disables the sending of ping packets if a valid address was previously defined. The default is to leave this field blank.
Retry Timeout (*retry_time*)	This field is optional. It controls how often the Link Aggregation device sends out a ping packet to poll the current adapter for link status. This is valid only when the Link Aggregation device has one or more adapters, a backup adapter is defined, and the **Internet Address to Ping** field contains a nonzero address. Specify the timeout value in seconds. The range of valid values is 1 - 100 seconds. The default value is 1 second.
Number of Retries (*num_retries*)	This field is optional. It specifies the number of lost ping packets before the Link Aggregation device switches adapters. This is valid only when the Link Aggregation device has one or more adapters, a backup adapter is defined, and the **Internet Address to Ping** field contains a nonzero address. The range of valid values is 2 - 100 retries. The default value is 3.
Enable Gigabit Ethernet Jumbo Frames (*use_jumbo_frame*)	This field is optional. To use this attribute, all of the underlying adapters, as well as the switch, must support jumbo frames. This works only with a Standard Ethernet (en) interface, not an IEEE 802.3 (et) interface.
Enable Alternate Address (*use_alt_addr*)	This field is optional. If you set this to yes, you can specify a MAC address that you want the Link Aggregation device to use. If you set this option to no, the Link Aggregation device uses the MAC address of the first adapter.
Alternate Address (*alt_addr*)	If **Enable Alternate Address** is set to `yes`, specify the MAC address that you want to use. The address that you specify must start with `0x` and be a 12-digit hexadecimal address.

VLAN attributes

You can modify the following VLAN attributes.

Attribute	Value
VLAN Tag ID (*vlan_tag_id*)	The unique ID associated with the VLAN driver. You can specify in the range 1 - 4094.
Base Adapter (*base_adapter*)	The network adapter to which the VLAN device driver is connected.

Shared Ethernet Adapter QoS attribute

You can modify the following qos_mode attribute.

disabled mode
This is the default mode. VLAN traffic is not inspected for the priority field. For example,

```
chdev -dev <sea device name> -attr qos_mode=disabled
```

strict mode
More important traffic is bridged over less important traffic. This mode provides better performance and more bandwidth to more important traffic; however, it can result in substantial delays for less important traffic. For example,

```
chdev -dev <sea device name> -attr qos_mode=strict
```

loose mode
A cap is placed on each priority level, so that after a number of bytes are sent for each priority level, the next level is serviced. This method ensures that all packets will eventually be sent. More important traffic is given less bandwidth with this mode than with strict mode; however, the caps in loose mode are such that more bytes are sent for the more important traffic, so it still gets more bandwidth than less important traffic. For example,

```
chdev -dev <sea device name> -attr qos_mode=loose
```

Client-specific Shared Ethernet Adapter statistics

To gather network statistics at a client level, enable advanced accounting on the Shared Ethernet Adapter to provide more information about its network traffic. To enable client statistics, set the Shared Ethernet Adapter accounting attribute to enabled (the default value is disabled). When advanced accounting is enabled, the Shared Ethernet Adapter keeps track of the hardware (MAC) addresses of all of the packets it receives from the LPAR clients, and increments packet and byte counts for each client independently. After advanced accounting is enabled on the Shared Ethernet Adapter, you can generate a report to view per-client statistics by running the **seastat** command. The command must be run on the Shared Ethernet Adapter, which is actively bridging the traffic.

Note: Advanced accounting must be enabled on the Shared Ethernet Adapter before you can use the **seastat** command to print any statistics.

To enable advanced accounting on the Shared Ethernet Adapter, enter the following command:

```
chdev -dev <sea device name> -attr accounting=enabled
```

The following command displays per-client Shared Ethernet Adapter statistics. The optional -n flag disables name resolution on IP addresses.

```
seastat -d <sea device name> [-n]
```

The following command clears all of the per-client Shared Ethernet Adapter statistics that have been gathered:

```
seastat -d <sea device name> -c
```

Shared Ethernet Adapter failover statistics

Learn about Shared Ethernet Adapter failover statistics, such as high availability information and packet types, and view examples.

Statistic descriptions

Table 65. Descriptions of Shared Ethernet Adapter failover statistics

Statistic	Description
High availability	**Control Channel PVID** Port VLAN ID of the virtual Ethernet adapter that is used as the control channel. **Control Packets in** Number of packets that are received on the control channel. **Control Packets out** Number of packets that are sent on the control channel.
Packet types	**Keep-Alive Packets** Number of keep-alive packets that are received on the control channel. Keep-alive packets are received on the backup Shared Ethernet Adapter while the primary Shared Ethernet Adapter is active. **Recovery Packets** Number of recovery packets that are received on the control channel. Recovery packets are sent by the primary Shared Ethernet Adapter when it recovers from a failure and is ready to be active again. **Notify Packets** Number of notify packets that are received on the control channel. Notify packets are sent by the backup Shared Ethernet Adapter when it detects that the primary Shared Ethernet Adapter has recovered. **Limbo Packets** Number of limbo packets that are received on the control channel. Limbo packets are sent by the primary Shared Ethernet Adapter when it detects that its physical network is not operational, or when it cannot ping the specified remote host (to inform the backup that it needs to become active).

Table 65. Descriptions of Shared Ethernet Adapter failover statistics (continued)

Statistic	Description
State	The current state of the Shared Ethernet Adapter. **INIT** The Shared Ethernet Adapter failover protocol has just been initiated. **PRIMARY** The Shared Ethernet Adapter is actively connecting traffic between the VLANs to the network. **BACKUP** The Shared Ethernet Adapter is idle and not connecting traffic between the VLANs and the network. **PRIMARY_SH** The Shared Ethernet Adapter is configured in load sharing mode, and it connects traffic between a subset of VLANs and the network. **BACKUP_SH** The Shared Ethernet Adapter is configured in load sharing mode, and it connects traffic between a subset of VLANs that are not bridged by the primary Shared Ethernet Adapter. **RECOVERY** The primary Shared Ethernet Adapter recovered from a failure and is ready to be active again. **NOTIFY** The backup Shared Ethernet Adapter detected that the primary Shared Ethernet Adapter recovered from a failure and that it needs to become idle again. **LIMBO** One of the following situations is true: • The physical network is not operational. • The physical network's state is unknown. • The Shared Ethernet Adapter cannot ping the specified remote host.

Table 65. Descriptions of Shared Ethernet Adapter failover statistics (continued)

Statistic	Description
Bridge Mode	Describes to what level, if any, the Shared Ethernet Adapter is bridging traffic. **Unicast** The Shared Ethernet Adapter is sending and receiving only unicast traffic (no multicast or broadcast traffic). To avoid broadcast storms, the Shared Ethernet Adapter sends and receives unicast traffic only while it is in the INIT or the RECOVERY states. **All** The Shared Ethernet Adapter is sending and receiving all types of network traffic. **Partial** Used when the Shared Ethernet Adapter is in load sharing state (PRIMARY_SH or BACKUP_SH). In this mode, the Shared Ethernet Adapter bridges all types of traffic (unicast, broadcast, or multicast), but only for a subset of VLANs determined during load sharing negotiation. **None** The Shared Ethernet Adapter is not sending or receiving any network traffic.
Number of Times Server became Backup	Number of times the Shared Ethernet Adapter was active and became idle because of a failure.
Number of Times Server became Primary	Number of times the Shared Ethernet Adapter was idle and became active because the primary Shared Ethernet Adapter failed.

Table 65. Descriptions of Shared Ethernet Adapter failover statistics (continued)

Statistic	Description
High Availability Mode	How the Shared Ethernet Adapter behaves regarding the Shared Ethernet Adapter failover protocol. **Auto** The Shared Ethernet Adapter failover protocol determines whether the Shared Ethernet Adapter acts as the primary Shared Ethernet Adapter or as the backup Shared Ethernet Adapter. **Standby** The Shared Ethernet Adapter operates as a backup if there is another Shared Ethernet Adapter available to act as the primary. *Standby* causes a primary Shared Ethernet Adapter to become a backup Shared Ethernet Adapter if there is another Shared Ethernet Adapter that can become the primary Shared Ethernet Adapter. **Sharing** Sharing causes the backup Shared Ethernet Adapter to initiate a request for load sharing. The primary Shared Ethernet Adapter approves the request. After negotiation, both Shared Ethernet Adapters bridge traffic for an exclusive subset of VLANs. The **High Availability Mode** option must be set to *Sharing* on both Shared Ethernet Adapters, starting with the primary Shared Ethernet Adapter. **Priority** Specifies the trunk priority of the virtual Ethernet adapters of the Shared Ethernet Adapter. It is used by the Shared Ethernet Adapter protocol to determine which Shared Ethernet Adapter acts as the primary Shared Ethernet Adapter and which Shared Ethernet Adapter acts as the backup Shared Ethernet Adapter. Values range in the range 1 - 12, where a lesser number is favored to act as a primary Shared Ethernet Adapter.

Example statistics

Running the **entstat -all** command returns results similar to the following:

```
ETHERNET STATISTICS (ent8) :
Device Type: Shared Ethernet Adapter
Hardware Address: 00:0d:60:0c:05:00
Elapsed Time: 3 days 20 hours 34 minutes 26 seconds

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 7978002                              Packets: 5701362
Bytes: 919151749                              Bytes: 664049607
Interrupts: 3                                 Interrupts: 5523380
Transmit Errors: 0                            Receive Errors: 0
```

```
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0
Max Packets on S/W Transmit Queue: 2
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 1

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds
Broadcast Packets: 5312086                    Broadcast Packets: 3740225
Multicast Packets: 265589                     Multicast Packets: 194986
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 1

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 0
Driver Flags: Up Broadcast Running
    Simplex 64BitSupport ChecksumOffLoad
  DataRateSet

--------------------------------------------------------------
Statistics for adapters in the Shared Ethernet Adapter ent8
--------------------------------------------------------------
Number of adapters: 2
SEA Flags: 00000001
    < THREAD >
VLAN IDs :
    ent7: 1
Real Side Statistics:
    Packets received: 5701344
    Packets bridged: 5673198
    Packets consumed: 3963314
    Packets fragmented: 0
    Packets transmitted: 28685
    Packets dropped: 0
Virtual Side Statistics:
    Packets received: 0
    Packets bridged: 0
    Packets consumed: 0
    Packets fragmented: 0
    Packets transmitted: 5673253
    Packets dropped: 0
Other Statistics:
    Output packets generated: 28685
    Output packets dropped: 0
    Device output failures: 0
    Memory allocation failures: 0
    ICMP error packets sent: 0
    Non IP packets larger than MTU: 0
    Thread queue overflow packets: 0
High Availability Statistics:
    Control Channel PVID: 99
    Control Packets in: 0
    Control Packets out: 818825
Type of Packets Received:
    Keep-Alive Packets: 0
    Recovery Packets: 0
    Notify Packets: 0
    Limbo Packets: 0
    State: LIMBO
    Bridge Mode: All
    Number of Times Server became Backup: 0
    Number of Times Server became Primary: 0
    High Availability Mode: Auto
    Priority: 1

--------------------------------------------------------------
Real Adapter: ent2

ETHERNET STATISTICS (ent2) :
Device Type: 10/100 Mbps Ethernet PCI Adapter II (1410ff01)
Hardware Address: 00:0d:60:0c:05:00
```

```
Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 28684                                Packets: 5701362
Bytes: 3704108                                Bytes: 664049607
Interrupts: 3                                 Interrupts: 5523380
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 2
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 1

Broadcast Packets: 21                         Broadcast Packets: 3740225
Multicast Packets: 0                          Multicast Packets: 194986
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 1

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 200
Driver Flags: Up Broadcast Running
    Simplex Promiscuous AlternateAddress
    64BitSupport ChecksumOffload PrivateSegment LargeSend DataRateSet

10/100 Mbps Ethernet PCI Adapter II (1410ff01) Specific Statistics:
------------------------------------------------------------------------
Link Status: Up
Media Speed Selected: Auto negotiation
Media Speed Running: 100 Mbps Full Duplex
Receive Pool Buffer Size: 1024
No Receive Pool Buffer Errors: 0
Receive Buffer Too Small Errors: 0
Entries to transmit timeout routine: 0
Transmit IPsec packets: 0
Transmit IPsec packets dropped: 0
Receive IPsec packets: 0
Receive IPsec SA offload count: 0
Transmit Large Send packets: 0
Transmit Large Send packets dropped: 0
Packets with Transmit collisions:
  1 collisions: 0     6 collisions: 0     11 collisions: 0
  2 collisions: 0     7 collisions: 0     12 collisions: 0
  3 collisions: 0     8 collisions: 0     13 collisions: 0
  4 collisions: 0     9 collisions: 0     14 collisions: 0
  5 collisions: 0    10 collisions: 0     15 collisions: 0

--------------------------------------------------------------
Virtual Adapter: ent7

ETHERNET STATISTICS (ent7) :
Device Type: Virtual I/O Ethernet Adapter (l-lan)
Hardware Address: 8a:83:54:5b:4e:9a

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 7949318                              Packets: 0
Bytes: 915447641                              Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 5312065                    Broadcast Packets: 0
Multicast Packets: 265589                     Multicast Packets: 0
No Carrier Sense: 0                           CRC Errors: 0
```

```
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 20000
Driver Flags: Up Broadcast Running
    Simplex Promiscuous AllMulticast
    64BitSupport ChecksumOffload DataRateSet

Virtual I/O Ethernet Adapter (l-lan) Specific Statistics:
---------------------------------------------------------
RQ Lingth: 4481
No Copy Buffers: 0
Trunk Adapter: True
  Priority: 1  Active: True
Filter MCast Mode: False
Filters: 255
  Enabled: 1  Queued: 0  Overflow: 0
LAN State: Operational

Hypervisor Send Failures: 2371664
  Receiver Failures: 2371664
  Send Errors: 0

Hypervisor Receive Failures: 0

ILLAN Attributes: 0000000000003103 [0000000000003103]

PVID:  1      VIDs: None

Switch ID: ETHERNET0

Buffers   Reg  Alloc  Min  Max   MaxA  LowReg
 tiny     512   512   512  2048  512   512
 small    512   512   512  2048  512   512
 medium   128   128   128  256   128   128
 large    24    24    24   64    24    24
 huge     24    24    24   64    24    24

--------------------------------------------------------------
Control Adapter: ent9

ETHERNET STATISTICS (ent9) :
Device Type: Virtual I/O Ethernet Adapter (l-lan)
Hardware Address: 8a:83:54:5b:4e:9b

Transmit Statistics:                          Receive Statistics:
--------------------                          -------------------
Packets: 821297                               Packets: 0
Bytes: 21353722                               Bytes: 0
Interrupts: 0                                 Interrupts: 0
Transmit Errors: 0                            Receive Errors: 0
Packets Dropped: 0                            Packets Dropped: 0
                                              Bad Packets: 0

Max Packets on S/W Transmit Queue: 0
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Broadcast Packets: 821297                     Broadcast Packets: 0
Multicast Packets: 0                          Multicast Packets: 0
No Carrier Sense: 0                           CRC Errors: 0
DMA Underrun: 0                               DMA Overrun: 0
Lost CTS Errors: 0                            Alignment Errors: 0
Max Collision Errors: 0                       No Resource Errors: 0
Late Collision Errors: 0                      Receive Collision Errors: 0
Deferred: 0                                   Packet Too Short Errors: 0
SQE Test: 0                                   Packet Too Long Errors: 0
Timeout Errors: 0                             Packets Discarded by Adapter: 0
Single Collision Count: 0                     Receiver Start Count: 0
Multiple Collision Count: 0
```

```
Current HW Transmit Queue Length: 0

General Statistics:
-------------------
No mbuf Errors: 0
Adapter Reset Count: 0
Adapter Data Rate: 20000
Driver Flags: Up Broadcast Running
         Simplex 64BitSupport ChecksumOffload DataRateSet

Virtual I/O Ethernet Adapter (l-lan) Specific Statistics:
---------------------------------------------------------
RQ Length: 4481
No Copy Buffers: 0
Trunk Adapter: False
Filter MCast Mode: False
Filters: 255
  Enabled: 0  Queued: 0  Overflow: 0
LAN State: Operational

Hypervisor Send Failures: 0
  Receiver Failures: 0
  Send Errors: 0

Hypervisor Receive Failures: 0

ILLAN Attributes: 0000000000003002 [0000000000003002]

PVID:  99     VIDs:  None

Switch ID: ETHERNET0

Buffers        Reg  Alloc    Min     Max   MaxA  LowReg
 tiny          512    512    512    2048    512     512
 small         512    512    512    2048    512     512
 medium        128    128    128     256    128     128
 large          24     24     24      64     24      24
 huge           24     24     24      64     24      24
```

Shared Ethernet Adapter statistics

Learn about general Shared Ethernet Adapter statistics, such as VLAN IDs and packet information, and view examples.

Statistic descriptions

Table 66. Descriptions of Shared Ethernet Adapter statistics

Statistic	Description
Number of adapters	Includes the real adapter and all of the virtual adapters. **Note:** If you are using Shared Ethernet Adapter failover, then the control channel adapter is not included.

Table 66. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Shared Ethernet Adapter flags	Denotes the features that the Shared Ethernet Adapter is currently running. **THREAD** The Shared Ethernet Adapter is operating in threaded mode, where incoming packets are queued and processed by different threads; its absence denotes interrupt mode, where packets are processed in the same interrupt where they are received. **LARGESEND** The large send feature has been enabled on the Shared Ethernet Adapter. **JUMBO_FRAMES** The jumbo frames feature has been enabled on the Shared Ethernet Adapter. **GVRP** The GVRP feature has been enabled on the Shared Ethernet Adapter.
VLAN IDs	List of VLAN IDs that have access to the network through the Shared Ethernet Adapter (this includes PVID and all tagged VLANs).

Table 66. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Real adapters	**Packets received** Number of packets that are received on the physical network. **Packets bridged** Number of packets that are received on the physical network that were sent to the virtual network. **Packets consumed** Number of packets that are received on the physical network that were addressed to the interface configured over the Shared Ethernet Adapter. **Packets fragmented** Number of packets that are received on the physical network that were fragmented before being sent to the virtual network. They were fragmented because they were bigger than the outgoing adapter's Maximum Transmission Unit (MTU). **Packets transmitted** Number of packets that are sent on the physical network. This includes packets that are sent from the interface configured over the Shared Ethernet Adapter, as well as each packet sent from the virtual network to the physical network (including fragments). **Packets dropped** Number of packets that are received on the physical network that were dropped for one of the following reasons: • The packet was the oldest packet on a thread's queue and there was no space to accommodate a newly received packet. • The packet had an invalid VLAN ID and could not be processed. • The packet was addressed to the Shared Ethernet Adapter interface, but its interface had no filters registered. **Packets filtered (VLAN ID)** Number of packets that are received on the physical network that were not sent to the virtual network because of an unknown VLAN ID. **Packets filtered (Reserved address)** Number of packets that are received on the physical network that were not bridged to any of the trunk virtual Ethernet adapters because the destination MAC address is a reserved multicast address that is only useful for bridges.

Table 66. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Virtual adapters	**Packets received** Number of packets that are received on the virtual network. In other words, the number of packets received on all of the virtual adapters. **Packets bridged** Number of packets that are received on the virtual network that were sent to the physical network. **Packets consumed** Number of packets that are received on the virtual network that were addressed to the interface configured over the Shared Ethernet Adapter. **Packets fragmented** Number of packets that are received on the virtual network that were fragmented before being sent to the physical network. They were fragmented because they were bigger than the outgoing adapter's MTU. **Packets transmitted** Number of packets that are sent on the virtual network. This includes packets that are sent from the interface configured over the Shared Ethernet Adapter, as well as each packet sent from the physical network to the virtual network (including fragments). **Packets dropped** Number of packets that are received on the virtual network that were dropped for one of the following reasons: • The packet was the oldest packet on a thread's queue and there was no space to accommodate a newly received packet. • The packet was addressed to the Shared Ethernet Adapter interface, but its interface had no filters registered. **Packets filtered (VLAN ID)** In a shared high availability mode, the number of packets that are received on the virtual network and that were not sent to the physical network because they did not belong to the VLAN that is bridged by the shared Ethernet adapter.
Output packets generated	Number of packets with a valid VLAN tag or no VLAN tag sent out of the interface configured over the Shared Ethernet Adapter.

Table 66. Descriptions of Shared Ethernet Adapter statistics (continued)

Statistic	Description
Output packets dropped	Number of packets that are sent out of the interface configured over the Shared Ethernet Adapter that are dropped because of an invalid VLAN tag.
Device output failures	Number of packets that could not be sent due to underlying device errors. This includes errors that are sent on the physical network and virtual network, including fragments and Internet Control Message Protocol (ICMP) error packets generated by the Shared Ethernet Adapter.
Memory allocation failures	Number of packets that could not be sent because there was insufficient network memory to complete an operation.
ICMP error packets sent	Number of ICMP error packets that are successfully sent when a big packet could not be fragmented because the *don't fragment* bit was set.
Non IP packets larger than MTU	Number of packets that could not be sent because they were bigger than the outgoing adapter's MTU and could not be fragmented because they were not IP packets.
Thread queue overflow packets	Number of packets that were dropped from the thread queues because there was no space to accommodate a newly received packet.

The transmit statistic column indicates the sum of transmitted statistics for all SEAs. The receive statistic column indicates the sum of received statistics for all SEAs. For example, consider the following setup, where a Shared Ethernet Adapter has a real and a virtual adapter:

- ent5 = SEA
- ent0 = Real adapter
- ent1 = Virtual adapter

If a VIOClient receives 100 MB of data from a server, the SEA's real adapter records 100 MB on its receive statistic, and the SEA's virtual adapter records 100 MB on its transmit statistic. In this setup, the SEAs record 100 MB for transmit statistics column and 100 MB for receive statistics column.

If a VIOClient sends 300 MB of data to a server, the SEA's real adapter records 300 MB on its transmit statistic, and the SEA's virtual adapter records 300 MB on its receive statistic. In this setup, the SEAs record 300 MB for transmit statistics column and 300 MB for receive statistics column.

In a threaded mode, a section follows the statistics for each queue of each thread that handled packets. There is one queue per thread if QoS is disabled and seven queues per thread if QoS is enabled. Up to eight queues per thread are displayed if QoS mode is changed. You can use these statistics to verify whether the packets are distributed evenly between queues, whether the queues are sized correctly, and whether there are sufficient number of threads.

Table 67. Descriptions of Shared Ethernet Adapter per-queue statistics

Statistic	Description
Queue full dropped packets	Number of packets dropped from the thread queue because of lack of space to accommodate a newly received packet.

Table 67. Descriptions of Shared Ethernet Adapter per-queue statistics (continued)

Statistic	**Description**
Queue packets queued	Number of packets that are currently queued in the thread queue.
Queue average packets queued	Average number of packets present in the thread queue after a newly received packet is queued. A value of N indicates that on an average, there were N-1 packets already present in the queue when a new packet was queued.
Queue packets count	Total number of packets that have passed through the thread queue.
Queue max packets queued	Maximum number of packets that are handled by the thread queue.

Example statistics

An example of the statistics for adapters in the Shared Ethernet Adapter is as follows:

```
-------------------------------------------------------------
Statistics for adapters in the Shared Ethernet Adapter ent5
-------------------------------------------------------------
Number of adapters: 3
SEA Flags: 00000001
    < THREAD >
VLAN Ids :
    ent3: 15
    ent2: 14 100 101
Real Side Statistics:
    Packets received: 10763329
    Packets bridged: 10718078
    Packets consumed: 10708048
    Packets fragmented: 0
    Packets transmitted: 181044
    Packets dropped: 0
    Packets filtered(VlanId): 0
    Packets filtered(Reserved address): 45243
Virtual Side Statistics:
    Packets received: 363027
    Packets bridged: 181044
    Packets consumed: 0
    Packets fragmented: 0
    Packets transmitted: 10900061
    Packets dropped: 0
    Packets filtered(VlanId): 0
Other Statistics:
    Output packets generated: 181983
    Output packets dropped: 0
    Device output failures: 0
    Memory allocation failures: 0
    ICMP error packets sent: 0
    Non IP packets larger than MTU: 0
    Thread queue overflow packets: 0


        SEA THREADS INFORMATION

        Thread .............. #0
    SEA Default Queue #8
    Queue full dropped packets: 0
    Queue packets queued: 0
    Queue average packets queued: 1
    Queue packets count: 1811500
    Queue max packets queued: 8

        Thread .............. #1
    SEA Default Queue #8
    Queue full dropped packets: 0
    Queue packets queued: 0
    Queue average packets queued: 1
    Queue packets count: 1105002
```

```
     Queue max packets queued: 15

         Thread .............. #2
     SEA Default Queue #8
     Queue full dropped packets: 0
     Queue packets queued: 0
     Queue average packets queued: 1
     Queue packets count: 2213623
     Queue max packets queued: 12

         Thread .............. #3
     SEA Default Queue #8
     Queue full dropped packets: 0
     Queue packets queued: 0
     Queue average packets queued: 1
     Queue packets count: 502088
     Queue max packets queued: 12

         Thread .............. #4
     SEA Default Queue #8
     Queue full dropped packets: 0
     Queue packets queued: 0
     Queue average packets queued: 1
     Queue packets count: 654478
     Queue max packets queued: 12

         Thread .............. #5
     SEA Default Queue #8
     Queue full dropped packets: 0
     Queue packets queued: 0
     Queue average packets queued: 1
     Queue packets count: 2735294
     Queue max packets queued: 12

         Thread .............. #6
     SEA Default Queue #8
     Queue full dropped packets: 0
     Queue packets queued: 0
     Queue average packets queued: 1
     Queue packets count: 2104371
     Queue max packets queued: 12
```

User types for the Virtual I/O Server

Learn about Virtual I/O Server user types and their user permissions.

The Virtual I/O Server has the following user types: prime administrator, system administrator, service representative user, and development engineer user. After installation, the only user type that is active is the prime administrator.

Prime administrator

The prime administrator (**padmin**) user ID is the only user ID that is enabled after installation of the Virtual I/O Server and can run every Virtual I/O Server command. There can be only one prime administrator in the Virtual I/O Server.

System administrator

The system administrator user ID has access to all commands except the following commands:

- **lsfailedlogin**
- **lsgcl**
- **mirrorios**
- **mkuser**
- **oem_setup_env**
- **rmuser**
- **shutdown**
- **unmirrorios**

The prime administrator can create an unlimited number of system administrator IDs.

Service representative

Create the service representative (SR) user so that an IBM service representative can log in to the system and perform diagnostic routines. Upon logging in, the SR user is placed directly into the diagnostic menus.

Development engineer

Create a Development engineer (DE) user ID so that an IBM development engineer can log in to the system and debug problems.

View

This role is a read-only role and can perform list-type (ls) functions only. Users with this role do not have the authority to change the system configuration and do not have write permission to their home directories.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled "Cookies, Web Beacons and Other Technologies".

Programming interface information

This VIOS publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM VIOS Version 3.1.2.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, JBoss, OpenShift, Fedora, Hibernate, Ansible, CloudForms, RHCA, RHCE, RHCSA, Ceph, and Gluster are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Logical partitioning

Note

Before using this information and the product it supports, read the information in "Notices" on page 229.

This edition applies to IBM® AIX® Version 7.2, to IBM AIX Version 7.1, to IBM AIX Version 6.1, to IBM i 7.4 (product number 5770-SS1), to IBM Virtual I/O Server Version 3.1.2 and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

Contents

Logical partitioning

You can set up, manage, and troubleshoot AIX, IBM i, Linux®, and Virtual I/O Server logical partitions by using the Hardware Management Console (HMC), or Virtual Partition Manager. By creating logical partitions, you can reduce the footprint of your data center by consolidating servers, and maximize the use of system resources by sharing resources across logical partitions.

What's new in Logical partitioning

Read about new or changed information in Logical partitioning since the previous update of this topic collection.

March 2021

The topic "Dynamic Platform Optimizer " on page 146 is updated with versions of IBM Power Systems servers on which the Dynamic Platform Optimizer (DPO) function is supported.

November 2020

The following topics are updated:

- "Enabling the platform keystore capability on a logical partition" on page 69
- "Configuration requirements and restrictions for the remote restart capability of a logical partition " on page 67
- "Validating the simplified remote restart operation of a logical partition " on page 85
- "Remotely restarting a logical partition " on page 86

May 2020

- The following topics are new or updated for single root I/O virtualization (SR-IOV) logical ports:
  - "Configuration requirements and restrictions for the remote restart capability of a logical partition " on page 67
  - "Adding a single root I/O virtualization logical port to a logical partition dynamically" on page 157
  - "Viewing migratable SR-IOV logical ports and SR-IOV backup virtual devices" on page 157
  - "Modifying a single root I/O virtualization logical port that is assigned to a logical partition dynamically" on page 157
  - "Removing a single root I/O virtualization logical port from a logical partition dynamically" on page 158
  - "Creating a profile with migratable single root I/O virtualization logical ports" on page 158
  - "Recovering a migratable single root I/O virtualization logical port" on page 158

October 2019

- The following topics are new or updated for single root I/O virtualization (SR-IOV) logical ports:
  - "Configuration requirements and restrictions for the remote restart capability of a logical partition " on page 67
  - "Adding a single root I/O virtualization logical port to a logical partition dynamically" on page 157
  - "Viewing migratable SR-IOV logical ports and SR-IOV backup virtual devices" on page 157
  - "Modifying a single root I/O virtualization logical port that is assigned to a logical partition dynamically" on page 157

- “Removing a single root I/O virtualization logical port from a logical partition dynamically” on page 158
- “Creating a profile with migratable single root I/O virtualization logical ports” on page 158
- “Recovering a migratable single root I/O virtualization logical port” on page 158

August 2018

- The following topic is new for the hardware accelerator:
  - “Enabling user mode access to the hardware accelerator” on page 202
- The following topic was updated for mapping the virtual I/O resources of a logical partition to the Virtual I/O Servers on the destination server:
  - “Viewing the details of a simplified remote restart operation” on page 84

Logical partition overview

Logical partitioning is the ability to make a server run as if it were two or more independent servers. When you logically partition a server, you divide the resources on the server into subsets called logical partitions. You can install software on a logical partition, and the logical partition runs as an independent logical server with the resources that you allocated to the logical partition.

You can assign processors, memory, and input/output devices to logical partitions. You can run AIX, IBM i, Linux, and the Virtual I/O Server in logical partitions. The Virtual I/O Server provides virtual I/O resources to other logical partitions with general-purpose operating systems.

Logical partitions share a few system attributes, such as the system serial number, system model, and processor feature code. All other system attributes can vary from one logical partition to another.

You can create a maximum of 1000 logical partitions on a server. You must use tools to create logical partitions on your servers. The tool that you use to create logical partitions on each server depends on the server model and the operating systems and features that you want to use on the server.

Hardware Management Console

The Hardware Management Console (HMC) is a hardware appliance that you can use to configure and control one or more managed systems. You can use the HMC to create and manage logical partitions and activate Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information to service and support for analysis.

The HMC also provides terminal emulation for the logical partitions on your managed system. You can connect to logical partitions from the HMC itself, or you can set up the HMC so that you can connect to logical partitions remotely through the HMC. HMC terminal emulation provides a dependable connection that you can use if no other terminal device is connected or operational. HMC terminal emulation is useful during initial system setup before you configure your terminal of choice.

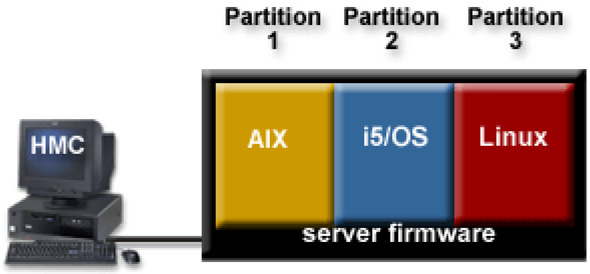

In this figure, you can see the logical partitions and the server firmware on the server. The server firmware is code that is stored in system flash memory on the server. The server firmware directly controls the resource allocations on the server and the communications between logical partitions on the server. The

HMC connects with the server firmware and specifies how the server firmware allocates resources to the managed system.

If you use a single HMC to manage a server, and the HMC malfunctions or becomes disconnected from the server firmware, then the server continues to run, but you cannot change the logical partition configuration of the server. If required, you can attach an extra HMC to act as a backup and to provide a redundant path between the server and service and support.

Partitioning by using the HMC is supported on all IBM Power Systems models, although some models require you to enter a PowerVM® Editions activation code before partitioning the managed system.

The PowerVM NovaLink architecture enables management of highly scalable cloud deployment by using the PowerVM technology and OpenStack solutions. The architecture provides a direct OpenStack connection to a PowerVM server. The NovaLink partition runs the Linux operating system and the partition runs on a server that is virtualized by PowerVM. The server is managed by PowerVC or other OpenStack solutions.

When a server is co-managed by the HMC and PowerVM NovaLink, and PowerVM NovaLink is in the master mode, you can run partition change operations only by using PowerVM NovaLink. If you want to run partition change operations by using the HMC, you must set the HMC to the master mode. Run the following command from the command line to set the HMC to the master mode:

```
chcomgmt -m <managed system> -o setmaster -t norm
```

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

A partition profile specifies the desired system resources for the logical partition and the minimum and maximum amounts of system resources that the logical partition can have. The system resources specified within a partition profile include processors, memory, and I/O resources. The partition profile can also specify certain operating settings for the logical partition. For example, you can set a partition profile such that, when the partition profile is activated, the logical partition is set to start automatically the next time that you power on the managed system.

Note: When a server is co-managed by the HMC, and PowerVM NovaLink, partition profiles are not supported.

Each logical partition on a managed system that is managed by an HMC has at least one partition profile. If you want, you can create more partition profiles with different resource specifications for your logical partition. If you create multiple partition profiles, you can designate any partition profile on the logical partition to be the default partition profile. The HMC activates the default profile if you do not select a specific partition profile to be activated. Only one partition profile can be active at one time. To activate another partition profile for a logical partition, you must shut down the logical partition before you activate the other partition profile.

A partition profile is identified by the logical partition ID and partition profile name. Logical partition IDs are whole numbers that are used to identify each logical partition that you create on a managed system, and partition profile names identify the partition profiles that you create for each logical partition. Each partition profile on a logical partition must have a unique partition profile name, but you can use a partition profile name for different logical partitions on a single managed system. For example, logical partition 1 cannot have more than one partition profile with a partition profile name of `normal`, but you can create a `normal` partition profile for each logical partition on the managed system.

When you create a partition profile, the HMC shows you all of the resources available on your system. The HMC does not verify whether another partition profile is currently using a portion of these resources. Therefore, it is possible for you to over commit resources. When you activate a logical partition by using a partition profile, the system attempts to start the logical partition by using the resources that are specified in the partition profile. If the minimum resources that are specified in the partition profile are not available on the managed system, the logical partition cannot be started by using the partition profile.

For example, you have four processors on your managed system. Logical partition 1 with partition profile A has three processors, and logical partition 2 with partition profile B has two processors. If you attempt to activate both of these partition profiles at the same time, logical partition 2 with partition profile B fails to activate because you over committed processor resources.

When you shut down a logical partition and reactivate the logical partition by using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition by using dynamic partitioning are lost when you reactivate the logical partition that uses a partition profile. This is preferred when you want to undo dynamic partitioning changes to the logical partition. However, this is not preferred if you want to reactivate the logical partition by using the resource specifications that the logical partition had when you shut down the managed system. Therefore, it is best to maintain your partition profiles up-to-date with the latest resource specifications. You can save the current configuration of the logical partition as a partition profile. Thereby, you do not need to change partition profiles manually.

If you shut down a logical partition whose partition profiles are not up-to-date, and if the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system by using the partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

You must activate a logical partition by activating a partition profile at least once. After that, you can activate the logical partition based on its current configuration data that is saved in the hypervisor. Logical partitions start faster when activated based on their current configuration data than when activated with a partition profile.

Processor resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of processor resources that you want for the logical partition.

The allocated value is the resource amount that the logical partition gets if you do not over commit the resource on the managed system. If the allocated amount of resources is available when you activate the partition profile, then the logical partition starts with the allocated amount of resources. However, if the allocated amount of resources is not available when you activate the partition profile, then the resources on your managed system are over committed. If the amount of resources that are available on the managed system is equal to or greater than the minimum amount of resources in the partition profile, then the logical partition starts with the available amount of resources. If the minimum amount of resources is not met, then the logical partition does not start.

If the managed system allows the configuration of multiple shared processor pools, then you can limit the number of processors that are used by a specific group of logical partitions by configuring a shared processor pool for those logical partitions and reassigning those logical partitions to that shared processor pool. For example, if you use per-processor licensing for IBM i, and you have a limited number of IBM i licenses for your managed system, you can create a shared processor pool for the IBM i logical partitions on the managed system and set the maximum number of processing units for that shared processor pool to be equal to the number of IBM i licenses on the managed system. If you configure a shared processor pool and assign logical partitions to that shared processor pool, the number of processing units that are used by those logical partitions plus the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool cannot exceed the maximum number of processing units that you set for that shared processor pool.

If you create a partition profile that is set to use shared processors, the HMC calculates a minimum, maximum, and allocated number of virtual processors for the partition profile. The calculation of virtual processors is based on the minimum, maximum, and allocated number of processing units that you specify for the partition profile. By default, the virtual processor settings are calculated as follows:

- The default minimum number of virtual processors is the minimum number of processing units (rounded up to the next whole number). For example, if the minimum number of processing units is 0.8, the default minimum number of virtual processors is 1.
- The default allocated number of virtual processors is the allocated number of processing units (rounded up to the next whole number). For example, if the allocated number of processing units is 2.8, the default allocated number of virtual processors is 3.
- The default maximum number of virtual processors is the maximum number of processing units rounded up to the next whole number and multiplied by two. For example, if the maximum number of processing units is 3.2, the default maximum number of virtual processors is 8 (four times 2).

When you activate the logical partition that uses the partition profile on the HMC, the allocated number of virtual processors is assigned to the logical partition. You can then use dynamic partitioning to change the number of virtual processors to any number between the minimum and maximum values, providede the number of virtual processors is greater than the number of processing units that are assigned to the logical partition. Before you change the default settings, performance modeling must be performed.

For example, you create a partition profile on the HMC with the following processor unit settings.

Minimum processing units	1.25
Allocated processing units	3.80
Maximum processing units	5.00

The default virtual processor settings for this partition profile on the HMC are as follows.

Minimum virtual processors	2
Allocated virtual processors	4
Maximum virtual processors	10

When you activate the logical partition by using this partition profile on the HMC, four processors are available to the operating system because the logical partition is activated with the allocated value of four virtual processors. Each of these virtual processors has 0.95 processing units to support the work that is assigned to the processor. After the logical partition is activated, you can use dynamic partitioning to change the number of virtual processors on the logical partition to any between number 2 - 10, provided the number of virtual processors is greater than the number of processing units that are assigned to the logical partition. If you increase the number of virtual processors, less processing power is available to support the work that is assigned to each processor.

Related concepts

Processors

A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

Memory resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of memory resources that you want for the logical partition.

When you create a partition profile that is set to use dedicated memory, the allocated, minimum, and maximum amounts of memory that you specify refer to physical memory in the system. If the allocated amount of physical memory is available on the managed system when you activate the partition profile, the logical partition starts with the allocated amount of physical memory. However, if the allocated amount of physical memory is not available when you activate the partition profile, the physical memory on your managed system is over committed. In that case, if the amount of physical memory that is available on the managed system is equal to or greater than the minimum amount of physical memory in the partition profile, the logical partition starts with the available amount of physical memory. If the minimum amount of physical memory is not available, then the logical partition does not start.

When you create a partition profile that is set to use shared memory, the allocated, minimum, and maximum amounts of memory that you specify refer to logical memory. When you activate the partition profile, the logical partition starts with the allocated amount of logical memory. You can dynamically add and remove logical memory to and from a running logical partition within the minimum and maximum values set in the partition profile.

Related concepts

Memory

Processors use memory to temporarily hold information. Memory requirements for logical partitions depend on the logical partition configuration, I/O resources assigned, and applications used.

I/O device assignment in partition profiles

I/O devices are assigned to partition profiles either on a slot-by-slot basis, or on logical port basis in the case of shared mode single root I/O virtualization (SR-IOV) adapters. For I/O devices that are assigned to partition profiles on a slot-by-slot basis, most I/O devices can be assigned to a partition profile on the HMC as required or as allocated. For SR-IOV logical ports, I/O devices are always assigned to a profile as required.

- If an I/O device is assigned to a partition profile as required, then the partition profile cannot be successfully activated if the I/O device is unavailable or is in use by another logical partition. Also, after the logical partition starts, you cannot use dynamic partitioning to remove the required I/O device from the running logical partition or move the required I/O device to another logical partition. This setting is suitable for devices that are required for the continuous operation of the logical partition (such as disk drives).
- If an I/O device is assigned to a partition profile as desired, then the partition profile can be successfully activated if the I/O device is unavailable or is in use by another logical partition. The desired I/O device can also be unconfigured in the operating system or system software and removed from the running logical partition or moved to another logical partition by using dynamic partitioning. This setting is suitable for devices that you want to share among multiple logical partitions (such as optical drives or tape drives).

The exception to this rule is host channel adapters (HCAs), which are added to partition profiles on the HMC as required. Each physical HCA contains a set of 64 globally unique IDs (GUIDs) that can be assigned to partition profiles. You can assign multiple GUIDs to each partition profile, but you can assign only one GUID from each physical HCA to each partition profile. Also, each GUID can be used by only one logical partition at a time. You can create multiple partition profiles with the same GUID, but only one of those partition profiles can be activated at a time.

You can change the required or desired setting within any partition profile for any I/O device at any time. Changes to the required or desired setting for an I/O device take effect immediately, even if the logical partition is running. For example, you want to move a tape device from one running logical partition to another, and the I/O device is required in the active partition profile for the source logical partition. You can access the active partition profile for the source logical partition, set the tape device to be allocated, and then unconfigure and move the tape device to the other logical partition without having to restart either logical partitions.

If you create an IBM i logical partition by using the HMC, you must tag I/O devices to perform certain functions for that IBM i logical partition.

Related concepts

Tagged resources for IBM i logical partitions

When you create an IBM i logical partition using the Hardware Management Console (HMC), you must tag I/O adapters (IOAs) to perform specific functions for the IBM i logical partition.

Related reference

Virtual adapters

With virtual adapters, you can connect logical partitions with each other without using physical hardware. Operating systems can display, configure, and use virtual adapters just like they can display, configure, and use physical adapters. Depending on the operating environment used by the logical partition, you can create virtual Ethernet adapters, virtual Fibre Channel adapters, virtual Small Computer Serial Interface (SCSI) adapters, and virtual serial adapters for a logical partition.

Partition profiles that use all of the system resources

You can create partition profiles on your HMC that specify all of the resources on the managed system. If you activate a logical partition by using such a partition profile, then the managed system assigns all of its resources to the logical partition.

If you add more resources to the managed system, the managed system automatically assigns the added resources to the logical partition when the profile is activated. The profile must be activated while the server is in the partition standby state because restarting the logical partition automatically does not assign newly added processor and memory resources. You do not have to change the partition profile for the managed system to assign the additional resources to the logical partition.

You cannot activate a logical partition that uses a partition profile that specifies all of the system resources if any other logical partition is running. However, after the logical partition is activated with all of the system resources, you can remove most processor and memory resources and all I/O resources from the logical partition by using dynamic partitioning. This allows you to start other logical partitions by using the resources that you remove from the logical partition. An implicit minimum amount of processor and memory resources is reserved for the logical partition that uses all of the system resources. Hence, you cannot remove all processor and memory resources from such a logical partition.

System profile

A system profile is an ordered list of partition profiles that is used by the Hardware Management Console (HMC) to start the logical partitions on a managed system in a specific configuration.

When you activate the system profile, the managed system attempts to activate each partition profile in the system profile in the order specified. A system profile helps you to activate or change the managed system from one complete set of logical partition configurations to another.

You can create a system profile and specify more resources to partition profiles than the resources that are available on the managed system. You can use the HMC to validate the system profile against the currently available system resources and against the total system resources. Validating your system profile ensures that your I/O devices and processing resources are not over committed, and it increases the likelihood that the system profile can be activated. The validation process estimates the amount of memory that is needed to activate all of the partition profiles in the system profile. A system profile can pass validation but might not have enough memory to be activated.

System profiles cannot include partition profiles that specify shared memory. In other words, logical partitions that use shared memory cannot be activated by using a system profile.

Benefits of logical partitioning

When you create logical partitions on your server, you can consolidate servers, share system resources, create mixed environments, and run integrated clusters.

The following scenarios illustrate the benefits of partitioning your server:

Consolidating servers

A logically partitioned server can reduce the number of servers that are required within an enterprise. You can consolidate several servers into a single logically partitioned system. This eliminates the need for, and expense of, more equipment.

Sharing resources

You can quickly and easily move hardware resources from one logical partition to another logical partition as needs change. Technologies such as the Micro-Partitioning® technology, allow for processor resources to be shared automatically among logical partitions that use a shared processor pool. Similarly, the PowerVM Active Memory Sharing technology allows for memory resources to be shared automatically among logical partitions that use the shared memory pool. Other technologies, such as dynamic partitioning, allow for resources to be manually moved to, from, and between running logical partitions without shutting down or restarting the logical partitions.

Maintaining independent servers

Dedicating a portion of the resources (disk storage unit, processors, memory, and I/O devices) to a logical partition achieves logical isolation of software. If configured correctly, logical partitions also have some hardware fault tolerance. Batch and 5250 online transaction processing (OLTP) workloads, which might not run together on a single machine, can be isolated and run efficiently in separate partitions.

Creating a mixed production and test environment

You can create a combined production and test environment on the same server. The production logical partition can run your main business applications, and the test logical partition is used to test software. A failure in a test logical partition, while not necessarily planned, does not disrupt normal business operations.

Merging production and test environments

Partitioning enables separate logical partitions to be allocated for production and test servers, eliminating the need to purchase more hardware and software. When testing completes, the resources that are allocated to the test logical partition can be returned to the production logical partition or elsewhere. As new projects are developed, they can be built and tested on the same hardware on which they are eventually deployed.

Running integrated clusters

Using high-availability application software, your partitioned server can run as an integrated cluster. You can use an integrated cluster to protect your server from most unscheduled failures within a logical partition.

Although there are many benefits of creating logical partitions, consider the following points before you choose to use logical partitions:

- Processor and memory failures might result in the failure of the entire server with all of its logical partitions. (The failure of a single I/O device affects only the logical partition to which the I/O device belongs.) To reduce the possibility of system failure, you can use the Advanced System Management Interface (ASMI) to set the server to unconfigure failing processors or memory modules automatically. After the server unconfigures the failing processor or memory module, the server continues running without using the unconfigured processor or memory module.
- Administering a consolidated system might be more difficult in some ways than administering multiple smaller systems, particularly if the resources in the consolidated system are used at a level close to their capacity. If you anticipate that you will use your server at a level close to its capacity, consider ordering a server model that is capable of Capacity on Demand (CoD).

Related information

Capacity on Demand

Sharing resources between logical partitions

Although each logical partition acts as an independent server, the logical partitions on a server can share some types of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed.

The following list illustrates some of the ways in which logical partitions can share resources. For some server models, the features that are mentioned in this list are options for which you must obtain and enter an activation code:

- The Micro-Partitioning technology (or shared processing) allows logical partitions to share the processors in shared processor pools. Each logical partition that uses shared processors is assigned a specific amount of processor power from its shared processor pool. By default, each logical partition is set such that the logical partition uses no more than its assigned processor power. Optionally, you can set a logical partition such that the logical partition can use processor power that is not being used by other logical partitions in its shared processor pool. If you set the logical partition such that it can use unused processor power, the amount of processor power that the logical partition can use is limited by the virtual processor settings of the logical partition and by the amount of unused processor power available in the shared processor pool that is used by the logical partition.
- Logical partitions can share the memory in the shared memory pool by using the PowerVM Active Memory Sharing technology (or shared memory). Instead of assigning a dedicated amount of physical

memory to each logical partition that uses shared memory (also referred to as *shared memory partitions*), the hypervisor constantly provides the physical memory from the shared memory pool to the shared memory partitions as needed. The hypervisor provides portions of the shared memory pool that are not currently being used by shared memory partitions to other shared memory partitions that need to use the memory. When a shared memory partition needs more memory than the current amount of unused memory in the shared memory pool, the hypervisor stores a portion of the memory that belongs to the shared memory partition in auxiliary storage. Access to the auxiliary storage is provided by a Virtual I/O Server logical partition. When the operating system attempts to access data that is located in the auxiliary storage, the hypervisor directs a Virtual I/O Server to retrieve the data from the auxiliary storage and write it to the shared memory pool, so that the operating system can access the data. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition hardware feature, which also includes the license for the Virtual I/O Server software. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

- Dynamic partitioning allows you to manually move resources to, from, and between running logical partitions without shutting down or restarting the logical partitions. This allows you to share devices that logical partitions use occasionally. For example, if the logical partitions on your server use an optical drive occasionally, you can assign a single optical drive to multiple logical partitions as a desired device. The optical drive would belong to only one logical partition at a time, but you can use dynamic partitioning to move the optical drive between logical partitions as needed. Dynamic partitioning is not supported on servers that are managed by using the Virtual Partition Manager.
- Virtual I/O allows logical partitions to access and use I/O resources on other logical partitions. For example, virtual Ethernet allows you to create a virtual LAN that connects the logical partitions on your server to each other. If one of the logical partitions on the server has a physical Ethernet adapter that is connected to an external network, you can configure the operating system of that logical partition to connect the virtual LAN with the physical Ethernet adapter. This allows the logical partitions on the server to share a physical Ethernet connection to an external network.
- A Host Ethernet Adapter (HEA), or Integrated Virtual Ethernet (IVE), allows multiple logical partitions on the same server to share a single physical Ethernet adapter. Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. The logical partitions can then access external networks through the HEA without using an Ethernet bridge on another logical partition.

  Note: HEA is not supported on POWER9 processor-based server.
- The single root I/O virtualization (SR-IOV) specification defines extensions to the PCI Express (PCIe) specification. SR-IOV allows virtualization of the physical ports of an adapter. Hence, the ports can be shared by multiple partitions that are running simultaneously. For example, a single physical Ethernet port appears as several separate physical devices.

Related concepts

Shared processors
Shared processors are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as the *Micro-Partitioning* technology.

Shared memory
You can configure your system such that multiple logical partitions share a pool of physical memory. A shared memory environment includes the shared memory pool, logical partitions that use the shared memory in the shared memory pool, logical memory, I/O entitled memory, at least one Virtual I/O Server logical partition, and paging space devices.

Managed systems

A managed system is a single physical server plus the resources that are connected to the physical server. The physical server and the connected resources are managed by the physical server as a single unit.

Connected resources can include expansion units, towers, and drawers, and storage area network (SAN) resources that are assigned to the server.

You can install a single operating system on a managed system and use the managed system as a single server. Alternately, you can use a partitioning tool, such as the Hardware Management Console (HMC) to create multiple logical partitions on the managed system. The partitioning tool manages the logical partitions on the managed system.

Manufacturing default configuration

The manufacturing default configuration is the initial single partition setup of the managed system as received from your service provider.

When your system is in the manufacturing default configuration, you can install an operating system on the managed system and use the managed system as a nonpartitioned server. In this state, you do not have to manage the system by using a Hardware Management Console (HMC).

If you choose to attach an HMC to a managed system that is in the manufacturing default configuration for reasons other than partitioning (such as to activate Capacity on Demand), all of the physical hardware resources on the system are automatically assigned to the logical partition. If you add new physical hardware resources to the managed system, the resources are automatically assigned to the logical partition. However, to use the newly added resources, you must dynamically add the resources to the logical partition or restart the logical partition. You do not have to make any partitioning changes on the server if you do not want to do so.

However, if you use the HMC to create, delete, change, copy, or activate any logical partitions or partition profiles on the managed system, the system is then in the partition mode. You must then use the HMC to manage the managed system. If the server has at least one IBM i logical partition, then you must also change the managed system properties on the HMC so that one of the IBM i logical partitions on the managed system is the service partition for the managed system. If a managed system is managed by using an HMC, and you want to return the managed system to a nonpartitioned state, or if you want to partition the managed system with the or the Virtual Partition Manager, then you must follow a special procedure to reset the server.

Managed systems that are partitioned by using the Virtual Partition Manager are not managed by an HMC. If a managed system is managed by using the Virtual Partition Manager, then you do not have to reset the server to return the managed system to a nonpartitioned state. Also, you do not have to reset the server if you want to start using the Virtual Partition Manager instead of using an HMC. To start using an HMC, back up the data on each logical partition, attach the HMC to the server, create the logical partitions, and restore the data to the storage assigned to each logical partition.

Related concepts

Logical partitioning tools
You must use tools to create logical partitions on your servers. The tool that you use to create logical partitions on each server depends upon the server model and the operating systems and features that you want to use on the server.

Logical partitioning tools

You must use tools to create logical partitions on your servers. The tool that you use to create logical partitions on each server depends upon the server model and the operating systems and features that you want to use on the server.

Virtual Partition Manager

The Virtual Partition Manager is a feature of IBM i that allows you to create and manage one host partition and up to four client logical partitions, running either Linux or IBM i. You can use the Virtual Partition

Manager to create logical partitions on a server that does not have a Hardware Management Console (HMC).

To use the Virtual Partition Manager, you must first install IBM i on a nonpartitioned server. After you install IBM i, you can initiate a console session on IBM i and use System Service Tools (SST) to create and configure IBM i or Linux logical partitions. IBM i controls the resource allocations of the logical partitions on the server.

When you use the Virtual Partition Manager to create logical partitions on a server, SST can be used to create and manage the logical partitions. IBM Navigator for i offers an improved interface and more features for these functions. The console session that you use to access SST can be initiated by using Operations Console LAN.

Related information

Virtual Partition Manager: A Guide to Planning and Implementation

Creating IBM i Client Partitions Using Virtual Partition Manager

Physical and virtual hardware resources

When you create logical partitions on a managed system, you can assign the physical resources on the managed system directly to logical partitions. You can also share hardware resources among logical partitions by virtualizing those hardware resources. The methods that are used to virtualize and share hardware resources depend on the type of resource that you are sharing.

Processors

A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

You can set a logical partition to use either processors that are dedicated to the logical partition or processors that are shared with other logical partitions. If a logical partition uses dedicated processors, then you must assign processors (in increments of whole numbers) to the logical partition. A logical partition that uses dedicated processors cannot use any processing capacity beyond the processors that are assigned to the logical partition.

By default, all physical processors that are not dedicated to specific logical partitions are grouped together in a *shared processor pool*. You can assign a specific amount of the processing capacity in this shared processor pool to each logical partition that uses shared processors. Some models allow you to use the HMC to configure multiple shared processor pools. These models have a *default shared processor pool* that contains all the processor resources that do not belong to logical partitions that use dedicated processors or logical partitions that use other shared processor pools. The other shared processor pools on these models can be configured with a maximum processing unit value and a reserved processing unit value. The maximum processing unit value limits the total number of processors that can be used by the logical partitions in the shared processor pool. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool.

You can set a logical partition that uses shared processors to use as little as 0.10 processing units, which is approximately a 10th of the processing capacity of a single processor. When the firmware is at level 7.6, or later, you can set a logical partition that uses shared processors to use as little as 0.05 processing units, which is approximately a 20th of the processing capacity of a single processor. You can specify the number of processing units to be used by a shared processor logical partition down to the 100th of a processing unit. Also, you can set a shared processor logical partition such that, if the logical partition requires more processing capacity than its assigned number of processing units, the logical partition can use processor resources that are not assigned to any logical partition or processor resources that are assigned to another logical partition but that are not being used by the other logical partition. (Some server models might require you to enter an activation code before you can create logical partitions that use shared processors.)

You can assign up to the entire processing capacity on the managed system to a single logical partition, if the operating system and server model supports doing so. You can configure your managed system such that it does not comply with the software license agreement for your managed system, but you will receive out-of-compliance messages if you operate the managed system in such a configuration.

Automatic redistribution of work when a processor fails

If the server firmware detects that a processor is about to fail, or if a processor fails when the processor is not in use, then the server firmware creates a serviceable event. The server firmware can also unconfigure the failing processor automatically, depending upon the type of failure and the unconfiguration policies that you set up using the Advanced System Management Interface (ASMI). You can also unconfigure a failing processor manually using the ASMI.

When the server firmware unconfigures a failing processor, and if unassigned or unlicensed processors are not available on the managed system, the processor unconfiguration can cause the logical partition to which the processor is assigned to shut down. To avoid shutting down mission-critical workloads when your server firmware unconfigures a failing processor, you can use the HMC to set partition availablity priorities for the logical partitions on your managed system. A logical partition with a failing processor can acquire a replacement processor from one or more logical partitions with a lower partition-availability priority. The managed system can dynamically reduce the number of processors used by shared processor partitions with lower partition-availability priorities and use the freed processor resources to replace the failing processor. If this does not provide enough processor resources to replace the failing processor, the managed system can shut down logical partitions with lower partition-availability priorities and use those freed processor resources to replace the failing processor. The acquisition of a replacement processor allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

A logical partition can use processors only from logical partitions with lower partition-availability priorities. If all of the logical partitions on your managed system have the same partition-availability priority, then a logical partition can replace a failed processor only if the managed system has unlicensed or unassigned processors.

By default, the partition-availability priority of Virtual I/O Server logical partitions with virtual SCSI adapters is set to 191. The partition availablity priority of all other logical partitions is set to 127 by default.

Do not set the priority of Virtual I/O Server logical partitions to be lower than the priority of the logical partitions that use the resources on the Virtual I/O Server logical partition. Do not set the priority of IBM i logical partitions with virtual SCSI adapters to be lower than the priority of the logical partitions that use the resources on the IBM i logical partition. If the managed system shuts down a logical partition because of its partition availability priority, all logical partitions that use the resources on that logical partition are also shut down.

If a processor fails when the processor is in use, then the entire managed system shuts down. When a processor failure causes the entire managed system to shut down, the system unconfigures the processor and restarts. The managed system attempts to start the logical partitions that were running at the time of the processor failure with their minimum processor values, in partition-availability priority order, with the logical partition with the highest partition-availability priority being started first. If the managed system does not have enough processor resources to start all of the logical partitions with their minimum processor values, then the managed system starts as many logical partitions as it can with their minimum processor values. If there are any processor resources remaining after the managed system has started the logical partitions, then the managed system distributes any remaining processor resources to the running logical partitions in proportion to their desired processor values.

Related concepts

Software licensing for IBM licensed programs on logical partitions
If you use IBM licensed programs such as AIX and IBM i on a server with logical partitions, consider how many software licenses are required for your logical partition configuration. Careful consideration of your software might help minimize the number of software licenses that you must purchase.

Processor resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of processor resources that you want for the logical partition.

Related information

Setting deconfiguration policies

Deconfiguring hardware

Dedicated processors

Dedicated processors are whole processors that are assigned to a single logical partition.

If you choose to assign dedicated processors to a logical partition, you must assign at least one processor to that logical partition. Likewise, if you choose to remove processor resources from a dedicated logical partition, you must remove at least one processor from the logical partition.

On systems that are managed by a Hardware Management Console (HMC), dedicated processors are assigned to logical partitions using partition profiles.

By default, a powered-off logical partition using dedicated processors provides processors to uncapped logical partitions that use shared processors. If the uncapped logical partition needs additional processor resources, the uncapped logical partition can use the idle processors that belong to the powered-off dedicated logical partition, if the total number of processors used by the uncapped logical partition does not exceed the virtual processors assigned to the uncapped logical partition, and if the use of these idle processors does not cause the shared processor pool to exceed its maximum processing units. When you power on the dedicated logical partition while the uncapped logical partition is using the processors, the activated logical partition regains all of its processing resources. If you use the HMC, you can prevent dedicated processors from being used in the shared processor pool by disabling this function in the partition properties panels.

You can also set the properties of a logical partition that uses dedicated processors such that, unused processing cycles on those dedicated processors can be made available to uncapped logical partitions while the dedicated processor logical partition is running. You can change the processor sharing mode of the dedicated processor logical partition at any time, without having to shut down and restart the logical partition.

Related concepts

Partition profile
A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Shared processors

Shared processors are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as the *Micro-Partitioning* technology.

Note: For some models, the Micro-Partitioning technology is an option for which you must obtain and enter a PowerVM Editions activation code.

By default, all physical processors that are not dedicated to specific logical partitions are grouped together in a *shared processor pool*. You can assign a specific amount of the processing capacity in this shared processor pool to each logical partition that uses shared processors. Some models allow you to use the HMC to configure multiple shared processor pools. These models have a *default shared processor pool* that contains all the processors that do not belong to logical partitions that use dedicated processors or logical partitions that use other shared processor pools. The other shared processor pools on these models can be configured with a maximum processing unit value and a reserved processing unit value. The maximum processing unit value limits the total number of processing units that can be used by the logical partitions in the shared processor pool. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool.

You can assign partial processors to a logical partition that uses shared processors. *Processing units* are a unit of measure for shared processing power across one or more virtual processors. One shared processing unit on one virtual processor accomplishes approximately the same work as one dedicated processor.

The minimum number of processing units depends on the firmware level.

Table 1. Firmware level and processing units per virtual processor

Firmware level	**Minimum number of processing units per virtual processor**
FW740, or earlier	0.10
FW760, or later	0.05

Some server models allow logical partitions to use only a portion of the total active processors on the managed system. Hence, the full processing capacity of the managed system cannot be assigned to logical partitions. This is true for server models with one or two processors, where a large portion of processor resources is used as overhead.

When the firmware is at level FW760, or later, overall server performance can be impacted when too many virtual processors are configured on the managed system. You can verify the number of configured virtual processors by using the **lshwres** command from the HMC command line. An example of the output after running the **lshwres** command follows:

```
lshwres -m sysname -r proc --level sys -F
curr_sys_virtual_procs,max_recommended_sys_virtual_procs
4,240
```

where:

- `curr_sys_virtual_procs` indicates the current number of configured virtual processors.
- `max_recommended_sys_virtual_procs` indicates the recommended maximum number of configured virtual processors.

It is suggested that the number of configured virtual processors must not exceed the maximum number so that server performance is not affected.

The maximum number of active virtual processors for a shared processor partition is limited by a number of factors. On Power 795, Power 870, Power 880, Power 870C, Power 880C and Power 980 model servers, the firmware has a limit of 128 active shared virtual processors per partition. On all other models of POWER7, Power 8, and POWER9, the firmware has a limit of 64 active shared virtual processors per partition.

Note: The limits on the number of active virtual processors for a shared processor partition is applicable for the firmware, but different operating systems and different operating system versions might impose limits lower than the firmware limits.

On HMC-managed systems, shared processors are assigned to logical partitions that use partition profiles.

Logical partitions that use shared processors can have a sharing mode of capped or uncapped. An *uncapped logical partition* is a logical partition that can use more processor power than its assigned processing capacity. The amount of processing capacity that an uncapped logical partition can use is limited only by the number of virtual processors assigned to the logical partition or the maximum processing unit that is allowed by the shared processor pool that the logical partition uses. In contrast, a *capped logical partition* is a logical partition that cannot use more processor power than its assigned processing units.

For example, logical partitions 2 and 3 are uncapped logical partitions, and logical partition 4 is a capped logical partition. Logical partitions 2 and 3 are each assigned 3.00 processing units and four virtual processors. Logical partition 2 currently uses only 1.00 of its 3.00 processing units, but logical partition 3 currently has a workload demand that requires 4.00 processing units. Because logical partition 3 is uncapped and has four virtual processors, the server firmware automatically allows logical partition 3 to use 1.00 processing units from logical partition 2. This increases the processing power for logical partition 3 to 4.00 processing units. Soon afterward, logical partition 2 increases its workload demand to 3.00 processing units. The server firmware therefore automatically returns 1.00 processing units to logical partition 2 so that logical partition 2 can use its full, assigned processing capacity again. Logical partition 4 is assigned 2.00 processing units and three virtual processors, but currently has a workload demand that requires 3.00 processing units. Because logical partition 4 is capped, logical partition 4 cannot use any unused processing units from logical partitions 2 or 3. However, if the workload demand of logical partition 4 decreases below 2.00 processing units, logical partitions 2 and 3 might use any unused processing units from logical partition 4.

By default, logical partitions that use shared processors are capped logical partitions. You can set a logical partition to be an uncapped logical partition if you want the logical partition to use more processing power than its assigned amount.

Although an uncapped logical partition can use more processor power than its assigned processing capacity, the uncapped logical partition can never use more processing units than its assigned number of virtual processors. Also, the logical partitions that use a shared processor pool can never use more processing units than the maximum processing units configured for the shared processor pool.

If multiple uncapped logical partitions need more processor capacity at the same time, the server can distribute the unused processing capacity to all uncapped logical partitions. This distribution process is determined by the uncapped weight of each of the logical partitions.

Uncapped weight is a number in the range of 0 through 255 that you set for each uncapped logical partition in the shared processor pool. On the HMC, you can choose from any of the 256 possible uncapped weight values. By setting the uncapped weight (255 being the highest weight), any available unused capacity is distributed to contending logical partitions in proportion to the established value of the uncapped weight. The default uncapped weight value is 128. When you set the uncapped weight to 0, no unused capacity is distributed to the logical partition.

When the firmware is at level FW830, or earlier, uncapped weight is used only when more virtual processors consume unused resources than the available physical processors in the shared processor pool. If no contention exists for processor resources, the virtual processors are immediately distributed across the physical processors, independent of their uncapped weights. This can result in situations where the uncapped weights of the logical partitions do not exactly reflect the amount of unused capacity.

For example, logical partition 2 has one virtual processor and an uncapped weight of 100. Logical partition 3 also has one virtual processor, but an uncapped weight of 200. If logical partitions 2 and 3 both require more processing capacity, and there is not enough physical processor capacity to run both logical partitions, logical partition 3 receives two more processing units for every additional processing unit that logical partition 2 receives. If logical partitions 2 and 3 both require more processing capacity, and there is enough physical processor capacity to run both logical partitions, logical partition 2 and 3 receive an equal amount of unused capacity. In this situation, their uncapped weights are ignored.

When the firmware is at level FW840, or later, if multiple partitions are assigned to a shared processor pool, the uncapped weight is used as an indicator of how the processor resources must be distributed among the partitions in the shared processor pool with respect to the maximum amount of capacity that can be used by the shared processor pool. For example, logical partition 2 has one virtual processor and an uncapped weight of 100. Logical partition 3 also has one virtual processor, but an uncapped weight of 200. If logical partitions 2 and 3 both require more processing capacity, logical partition 3 receives two additional processing units for every additional processing unit that logical partition 2 receives.

The server distributes unused capacity among all of the uncapped shared processor partitions that are configured on the server, regardless of the shared processor pools to which they are assigned. For example, if you configure logical partition 1 to the default shared processor pool and you configure logical partitions 2 and 3 to a different shared processor pool, all three logical partitions compete for the same unused physical processor capacity in the server, even though they belong to different shared processor pools.

Related concepts

Sharing resources between logical partitions

Although each logical partition acts as an independent server, the logical partitions on a server can share some types of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed.

Partition profile
A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Virtual processors

A *virtual processor* is a representation of a physical processor core to the operating system of a logical partition that uses shared processors.

When you install and run an operating system on a server that is not partitioned, the operating system calculates the number of operations that it can perform concurrently by counting the number of processors on the server. For example, if you install an operating system on a server that has eight processors, and each processor can perform two operations at a time, the operating system can perform 16 operations at a time. In the same way, when you install and run an operating system on a logical partition that uses dedicated processors, the operating system calculates the number of operations that it can perform concurrently by counting the number of dedicated processors that are assigned to the logical partition. In both cases, the operating system can easily calculate how many operations it can perform at a time by counting the whole number of processors that are available to it.

However, when you install and run an operating system on a logical partition that uses shared processors, the operating system cannot calculate a whole number of operations from the fractional number of processing units that are assigned to the logical partition. The server firmware must therefore represent the processing power available to the operating system as a whole number of processors. This allows the operating system to calculate the number of concurrent operations that it can perform. A *virtual processor* is a representation of a physical processor to the operating system of a logical partition that uses shared processors.

The server firmware distributes processing units evenly among the virtual processors assigned to a logical partition. For example, if a logical partition has 1.80 processing units and two virtual processors, each virtual processor has 0.90 processing units to support its workload.

You can assign only a limited number of processing units for each virtual processor. The minimum number of processing units for each virtual processor is 0.10 (or ten virtual processors for every processing unit). When the firmware is at level FW760, or later, the minimum number of processing units is further lowered to 0.05 (or 20 virtual processors for every processing unit). The maximum number of processing units that can be assigned to each virtual processor is always 1.00. This means that a logical partition cannot use more processing units than the number of virtual processors that it is assigned, even if the logical partition is uncapped.

A logical partition generally performs best if the number of virtual processors is close to the number of processing units available to the logical partition. This lets the operating system manage to the workload on the logical partition effectively. In certain situations, you might be able to increase system performance slightly by increasing the number of virtual processors. If you increase the number of virtual processors, you increase the number of operations that can run concurrently. However, if you increase the number of virtual processors without increasing the number of processing units, the speed at which each operation runs will decrease. The operating system also cannot shift processing power between processes if the processing power is split between many virtual processors.

On HMC-managed systems, virtual processors are assigned to logical partitions using partition profiles.

Related concepts

Partition profile
A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile,

the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Software and firmware requirements for processing units

The minimum number of processing units of a logical partition depends on the firmware level and the version of the operating system that is running on the logical partition.

The following table lists the firmware levels and the operating system versions.

Table 2. Software and firmware requirements for processing units.

Minimum number of processing units per virtual processor	Firmware level	IBM i	AIX	Linux
0.10	FW740, or earlier	All	All	All
0.05	FW760, or later	All	Version 7 with Technology Level 2 or at Version 6 with Technology Level 8, or later.	A Linux distribution that supports the lower processor entitlement of 0.05 processing units per virtual processor

Memory

Processors use memory to temporarily hold information. Memory requirements for logical partitions depend on the logical partition configuration, I/O resources assigned, and applications used.

Memory can be assigned in increments of 16 MB, 32 MB, 64 MB, 128 MB, and 256 MB. The default memory block size varies according to the amount of configurable memory in the system.

Table 3. Default memory block size used for varying amounts of configurable memory

Amount of configurable memory	Default memory block size
Less than 4 GB	16 MB
Greater than 4 GB up to 8 GB	32 MB
Greater than 8 GB up to 16 GB	64 MB
Greater than 16 GB up to 32 GB	128 MB
Greater than 32 GB	256 MB

A logical partition can grow based on the amount of memory initially allocated to it. Memory is added and removed to and from logical partitions in units of logical memory blocks. For logical partitions that are initially sized less than 256 MB, the maximum size to which a logical partition can grow is 16 times its initial size (up to the assigned maximum memory of the logical partition). For logical partitions that are initially sized 256 MB or larger, the maximum size to which the logical partition can grow is 64 times its initial size (up to the assigned maximum memory of the logical partition). The smallest increment for adding or removing memory to or from a logical partition is 16 MB.

The memory block size can be changed by using the Logical Memory Block Size option in the Advanced System Management Interface (ASMI). The machine default value should only be changed under direction from your service provider. To change the memory block size, you must be a user with administrator authority, and you must shut down and restart the managed system for the change to take effect. If the minimum memory amount in any partition profile on the managed system is less than the new memory block size, you must also change the minimum memory amount in the partition profile.

Each logical partition has a hardware page table (HPT). The HPT ratio is the ratio of the HPT size to the maximum memory value for the logical partition. The HPT is allocated in the server firmware memory overhead for the logical partition, and the size of the HPT can affect the performance of the logical partition. The size of the HPT is determined by the following factors:

- The HPT ratio of 1/64 is the default value for IBM i logical partitions, and 1/128 for AIX and Linux logical partitions.

  Note: You can override the default value by using the HMC command-line interface to change the value in the partition profile.
- The maximum memory values that you establish for the logical partition (dedicated or shared)

On systems that are managed by a Hardware Management Console, memory is assigned to logical partitions using partition profiles.

Related concepts

Memory resource assignment in partition profiles
When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of memory resources that you want for the logical partition.

Dedicated memory

Dedicated memory is physical system memory that you assign to a logical partition that uses dedicated memory (hereafter referred to as a *dedicated memory partition*), and is reserved for use by the dedicated memory partition until you remove the memory from the dedicated memory partition or delete the dedicated memory partition.

Depending on the overall memory in your system and the maximum memory values you choose for each logical partition, the server firmware must have enough memory to perform logical partition tasks. The amount of memory required by the server firmware varies according to several factors. The following factors influence server firmware memory requirements:

- Number of dedicated memory partitions
- Partition environments of the dedicated memory partitions
- Number of physical and virtual I/O devices used by the dedicated memory partitions
- Maximum memory values assigned to the dedicated memory partitions

Note: Firmware level updates can also change the server firmware memory requirements. Larger memory block sizes can exaggerate the memory requirement change.

When selecting the maximum memory values for each dedicated memory partition, consider the following points:

- Maximum values affect the hardware page table (HPT) size for each dedicated memory partition
- The logical memory map size for each dedicated memory partition

If the server firmware detects that a memory module has failed or is about to fail, the server firmware creates a serviceable event. The server firmware can also unconfigure the failing memory module automatically, depending on the type of failure and the deconfiguration policies that you set up using the Advanced System Management Interface (ASMI). You can also unconfigure a failing memory module manually using the ASMI. If a memory module failure causes the entire managed system to shut down, the managed system restarts automatically if the managed system is in normal IPL mode. When the managed system restarts itself, or when you restart the managed system manually, the managed system attempts to start the dedicated memory partitions that were running at the time of the memory module failure with their minimum memory values. If the managed system does not have enough memory to start all of the dedicated memory partitions with their minimum memory values, the managed system starts as many dedicated memory partitions as it can with their minimum memory values. If any memory is remaining after the managed system has started as many dedicated memory partitions as it can, the managed system distributes the remaining memory resources to the running dedicated memory partitions in proportion to the required memory values.

Setting huge-page memory values for AIX dedicated memory partitions
Specify the number of 16 GB pages to allocate to an AIX huge-page memory pool.

About this task

On managed systems that support huge-page memory, you can use the Hardware Management Console (HMC) to set the value for the huge-page memory pool. You can also specify values for the number of huge pages to allocate to logical partitions.

Using huge pages can improve performance in specific environments that require a high degree of parallelism, such as in DB2® database. You can specify huge-page memory that can be used for the shared-memory buffer pools in DB2. For logically partitioned systems, you can specify the minimum, wanted, and maximum number of huge pages to assign to a logical partition when you create the logical partition or partition profile.

To set the huge-page memory values, the system must be in the powered-off state. The new value takes effect when you restart the system.

Calculating huge-page memory requirements for AIX dedicated memory partitions
Calculate the value for the number of pages to allocate to an AIX huge-page memory pool.

About this task

To use huge-page memory, you must ensure that your system has adequate memory resources to dedicate to the huge-page memory pool. The huge-page memory pool is a region of system memory that is mapped as 16 GB page segments and is managed separately from the base memory of the system. Before you can specify the value for huge-page memory, you must determine which applications you are running and what the huge-page requirements are for your applications.

Determining huge-page memory requirements for your application

The huge-page memory pool can be used to enhance performance for DB2 in AIX operating systems. To determine this value, calculate the amount of memory required for the shared buffer pool to support your DB2 applications. Refer to the DB2 recommendations for buffer pool memory for your particular application.

Note: The huge page memory allocation cannot be changed dynamically. When you change the number of huge pages on the server, the server must be rebooted. Changing the number of assigned huge pages for a logical partition requires that you restart the logical partition.

Considerations for calculating the huge-page values

The amount of huge-page memory that you can allocate is dependent on the following factors:

- Total amount of licensed memory for your server
- Amount of available memory after configured memory resources are accounted for
- Number of physical I/O connections to the server (each I/O connection requires memory for the I/O tables, which can be distributed among the physical memory regions and reduces the memory available for huge pages)
- Base memory configuration for logical partitions (huge pages are not calculated as part of the configured-partition memory allocation)
- The requirements that define a huge page, that is each huge page requires 16 GB of contiguous real memory and must start on a 16 GB memory boundary
- Huge pages cannot span processing units. Each processing unit requires 32 GB to ensure at least one 16 GB huge page when all of the other considerations previously listed are taken into account.

⚠ **Attention:** The server firmware reduces the huge-page pool size to satisfy some of these dependencies. When this occurs, error log entries are generated to indicate that the huge-page pool size was reduced. The error log reference code is B700 5300. The reference code details indicates hexadecimal values that indicate why the huge-page pool size could not be satisfied. The following example shows the possible entries and how to interpret the additional words in these entries:

- word 3 = 0x0000000100000106: This means that the huge-page pool was reduced to satisfy the system hardware configuration
  - word 4 = number of user-configured huge pages
  - word 5 = number of huge pages that could be provided
- word 3 = 0x0000000100000105: This means that the huge-page pool was reduced to satisfy the memory configuration of logical partitions
  - word 4 = number of huge pages before logical partitions were created
  - word 5 = firmware calculated number of huge pages after satisfying logical partition memory requirements
  - word 6 = number of huge pages in the pool

Calculating huge-page memory values

To calculate the server memory requirements to support huge pages, use the following steps:

Procedure

1. Determine the amount of base system memory and round that figure to the next 16 GB value.
2. Determine the number of I/O connection loops on your system and multiply the number by 16 GB. This calculation is required because the server needs a memory table for each I/O connection, and a 16 GB huge page cannot be located where an I/O table exists.
3. Take the larger of the values determined in step 1 and step 2. This is your base memory value.
4. Determine the number of huge pages that is required for your AIX applications. To determine this value, use the guidelines provided by your application documentation and the AIX Performance Management. Multiply the number of anticipated huge pages by 16 GB. Add this figure to the base figure determined in step 3. The resulting figure provides an estimate of the amount of licensed memory required to satisfy the logical partition and huge-page pool memory requirements for your system.

Shared memory

You can configure your system such that multiple logical partitions share a pool of physical memory. A shared memory environment includes the shared memory pool, logical partitions that use the shared memory in the shared memory pool, logical memory, I/O entitled memory, at least one Virtual I/O Server logical partition, and paging space devices.

Related concepts

Sharing resources between logical partitions

Although each logical partition acts as an independent server, the logical partitions on a server can share some types of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed.

Overview of shared memory

Shared memory is physical memory that is assigned to the shared memory pool and shared among multiple logical partitions. The *shared memory pool* is a defined collection of physical memory blocks that are managed as a single memory pool by the hypervisor. Logical partitions that you configure to use shared memory, share the memory in the pool with other shared memory partitions.

For example, you create a shared memory pool with 16 GB of physical memory. You then create three logical partitions, configure them to use shared memory, and activate the shared memory partitions. Each shared memory partition can use the 16 GB that are in the shared memory pool.

The hypervisor determines the amount of memory that is allocated from the shared memory pool to each shared memory partition based on the workload and memory configuration of each shared memory partition. When allocating the physical memory to the shared memory partitions, the hypervisor ensures that each shared memory partition can access only the memory that is allocated to the shared memory partition at any given time. A shared memory partition cannot access the physical memory that is allocated to another shared memory partition.

The amount of memory that you assign to the shared memory partitions can be greater than the amount of memory in the shared memory pool. For example, you can assign 12 GB to shared memory partition 1, 8 GB to shared memory partition 2, and 4 GB to shared memory partition 3. Together, the shared memory partitions use 24 GB of memory, but the shared memory pool has only 16 GB of memory. In this situation, the memory configuration is considered over committed.

Over committed memory configurations are possible because the hypervisor virtualizes and manages all of the memory for the shared memory partitions in the shared memory pool as follows:

1. When shared memory partitions are not actively using their memory pages, the hypervisor allocates those unused memory pages to shared memory partitions that currently need them. When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time. The hypervisor need not store any data in auxiliary storage.
2. When a shared memory partition requires more memory than the hypervisor can provide to it by allocating unused portions of the shared memory pool, the hypervisor stores some of the memory that belongs to a shared memory partition in the shared memory pool and stores the remainder of the memory that belongs to the shared memory partition in auxiliary storage. When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage. When the operating system attempts to access the data, the hypervisor might need to retrieve the data from auxiliary storage before the operating system can access it.

Because the memory that you assign to a shared memory partition might not always reside in the shared memory pool, the memory that you assign to a shared memory partition is *logical memory*. Logical memory is the address space assigned to a logical partition, that the operating system perceives as its main storage. For a shared memory partition, a subset of the logical memory is backed up by physical main storage (or physical memory from the shared memory pool) and the remaining logical memory is kept in auxiliary storage.

A Virtual I/O Server logical partition provides access to the auxiliary storage, or paging space devices, which are required for shared memory partitions in an over committed memory configuration. A *paging space device* is a physical or logical device that is used by a Virtual I/O Server to provide the paging space for a shared memory partition. The *paging space* is an area of nonvolatile storage that is used to hold portions of a shared memory partition's logical memory that does not reside in the shared memory pool. When the operating system that runs in a shared memory partition attempts to access data, and the data is located in the paging space device that is assigned to the shared memory partition, the hypervisor sends a request to a Virtual I/O Server to retrieve the data and write it to the shared memory pool so that the operating system can access it.

On systems that are managed by a Hardware Management Console (HMC), you can assign up to two Virtual I/O Server (VIOS) logical partitions to the shared memory pool at a time. When you assign two paging VIOS partitions to the shared memory pool, you can configure the paging space devices such that both paging VIOS partitions have access to the same paging space devices. When one paging VIOS partition becomes unavailable, the hypervisor sends a request to the other paging VIOS partition to retrieve the data on the paging space device.

You cannot configure paging VIOS partitions to use shared memory. Paging VIOS partitions do not use the memory in the shared memory pool. You assign paging VIOS partitions to the shared memory pool so that they can provide access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool.

Driven by workload demands from the shared memory partitions, the hypervisor manages over committed memory configurations by continually performing the following tasks:

- Allocating portions of physical memory from the shared memory pool to the shared memory partitions as needed.
- Requesting a paging VIOS partition to read and write data between the shared memory pool and the paging space devices as needed.

The ability to share memory among multiple logical partitions is known as the PowerVM Active Memory Sharing technology. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

Example: A shared memory configuration that is logically overcommitted
When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

The following figure shows a server with shared memory configuration that is logically over committed.

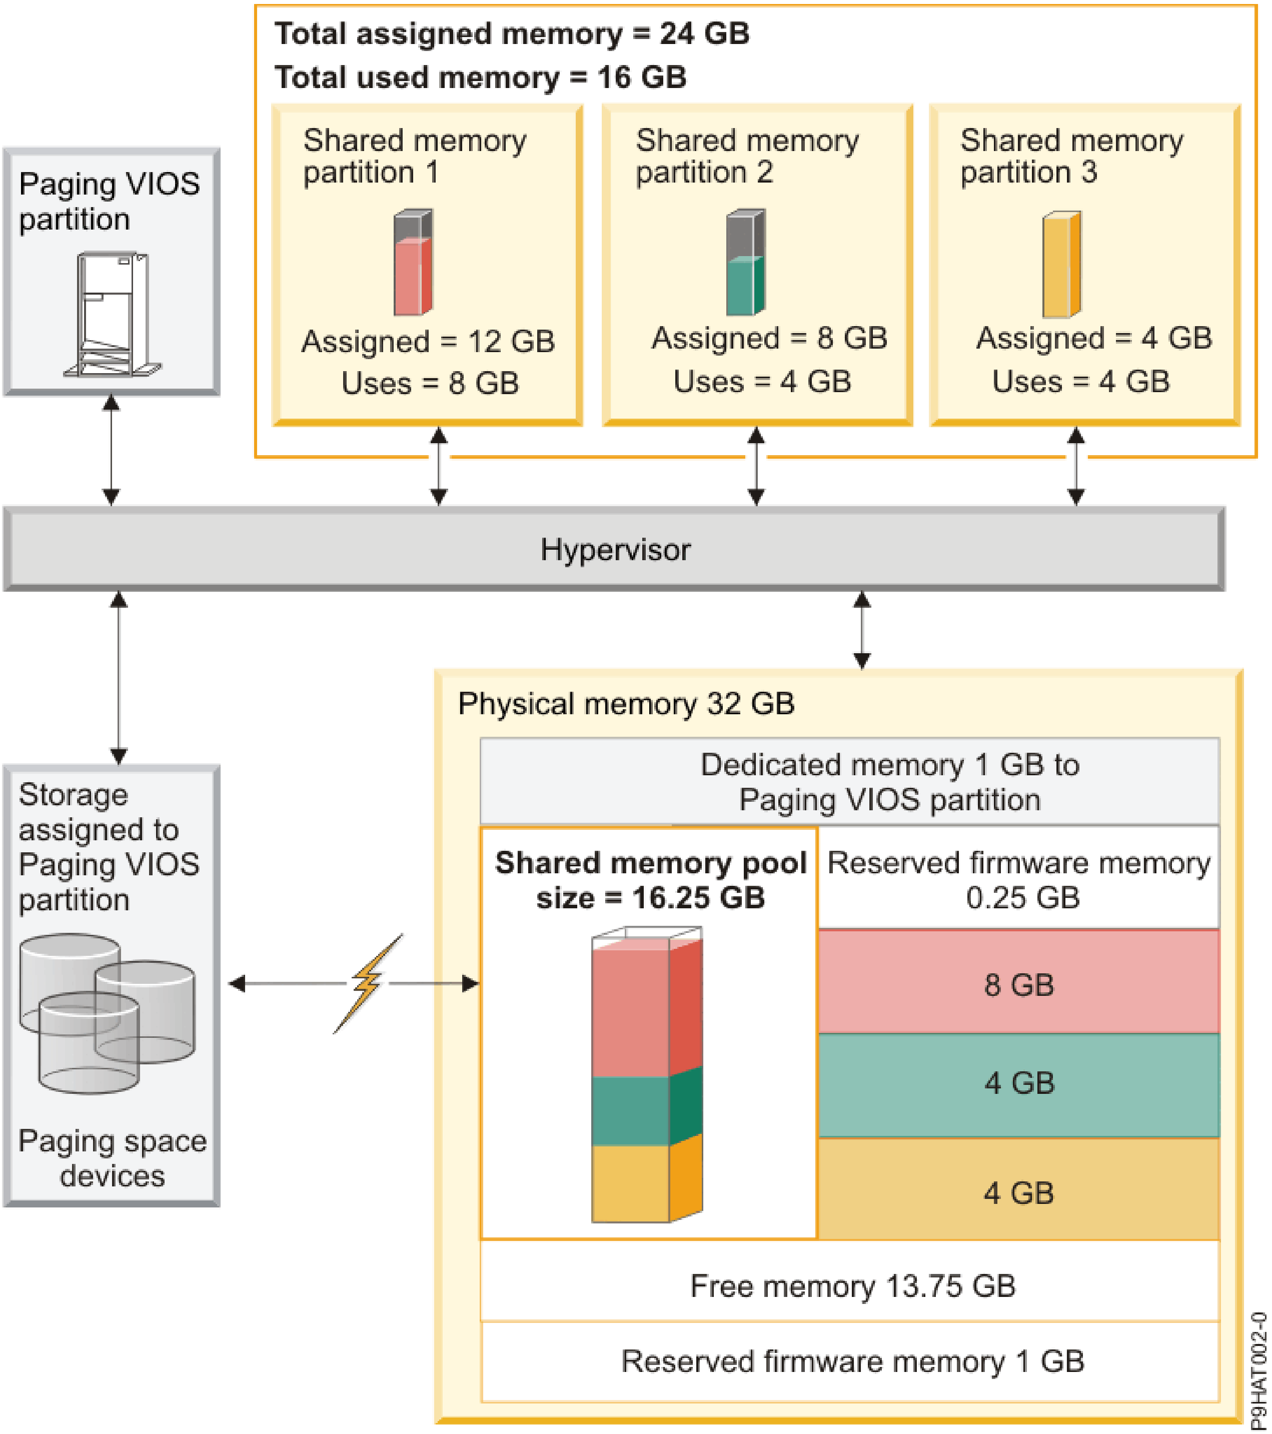
Total assigned memory = 24 GB
Total used memory = 16 GB
Shared memory partition 1
Assigned = 12 GB
Uses = 8 GB
Shared memory partition 2
Assigned = 8 GB
Uses = 4 GB
Shared memory partition 3
Assigned = 4 GB
Uses = 4 GB
Paging VIOS partition
Hypervisor
Physical memory 32 GB
Dedicated memory 1 GB to Paging VIOS partition
Shared memory pool size = 16.25 GB
Reserved firmware memory 0.25 GB
8 GB
4 GB
4 GB
Free memory 13.75 GB
Reserved firmware memory 1 GB
Storage assigned to Paging VIOS partition
Paging space devices
P9HAT002-0

Figure 1. A server with a shared memory configuration that is logically over committed

The figure shows a shared memory pool of 16.25 GB that is shared among three shared memory partitions. The hypervisor uses a small portion (0.25 GB) of the shared memory pool to manage the shared memory resources. The figure also shows one paging VIOS partition that owns all of the physical storage in the system. The physical storage contains a paging space device for each shared memory partition. The paging VIOS partition does not use the memory in the shared memory pool, but rather receives dedicated memory of 1 GB. Of the remaining system memory, 1 GB is reserved for the hypervisor so that it can manage other system resources, and 13.75 GB is free memory that is available for system growth. For example, you can dynamically add more memory to the shared memory pool or you can create additional dedicated memory partitions.

Shared memory partition 1 is assigned 12 GB of logical memory, Shared memory partition 2 is assigned 8 GB of logical memory, and Shared memory partition 3 is assigned 4 GB of logical memory. Together, the shared memory partitions are assigned 24 GB of logical memory, which is more than the 16.25 GB of logical memory allocated to the shared memory pool. Therefore, the memory configuration is over committed.

Shared memory partition 1 currently uses 8 GB of physical memory, Shared memory partition 2 currently uses 4 GB of physical memory, and Shared memory partition 3 currently uses 4 GB of physical memory. Together, the shared memory partitions currently use 16 GB of physical memory, which is equal to the amount of physical memory available to them in the shared memory pool. Therefore, the memory configuration is logically over committed. In other words, the shared memory pool contains enough physical memory for the hypervisor to allocate unused memory pages to shared memory partitions that need them. All of the memory currently used by the shared memory partitions resides in the shared memory pool.

Related concepts

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Example: A shared memory configuration that is physically overcommitted
When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

The following figure shows a server with shared memory configuration that is physically over committed.

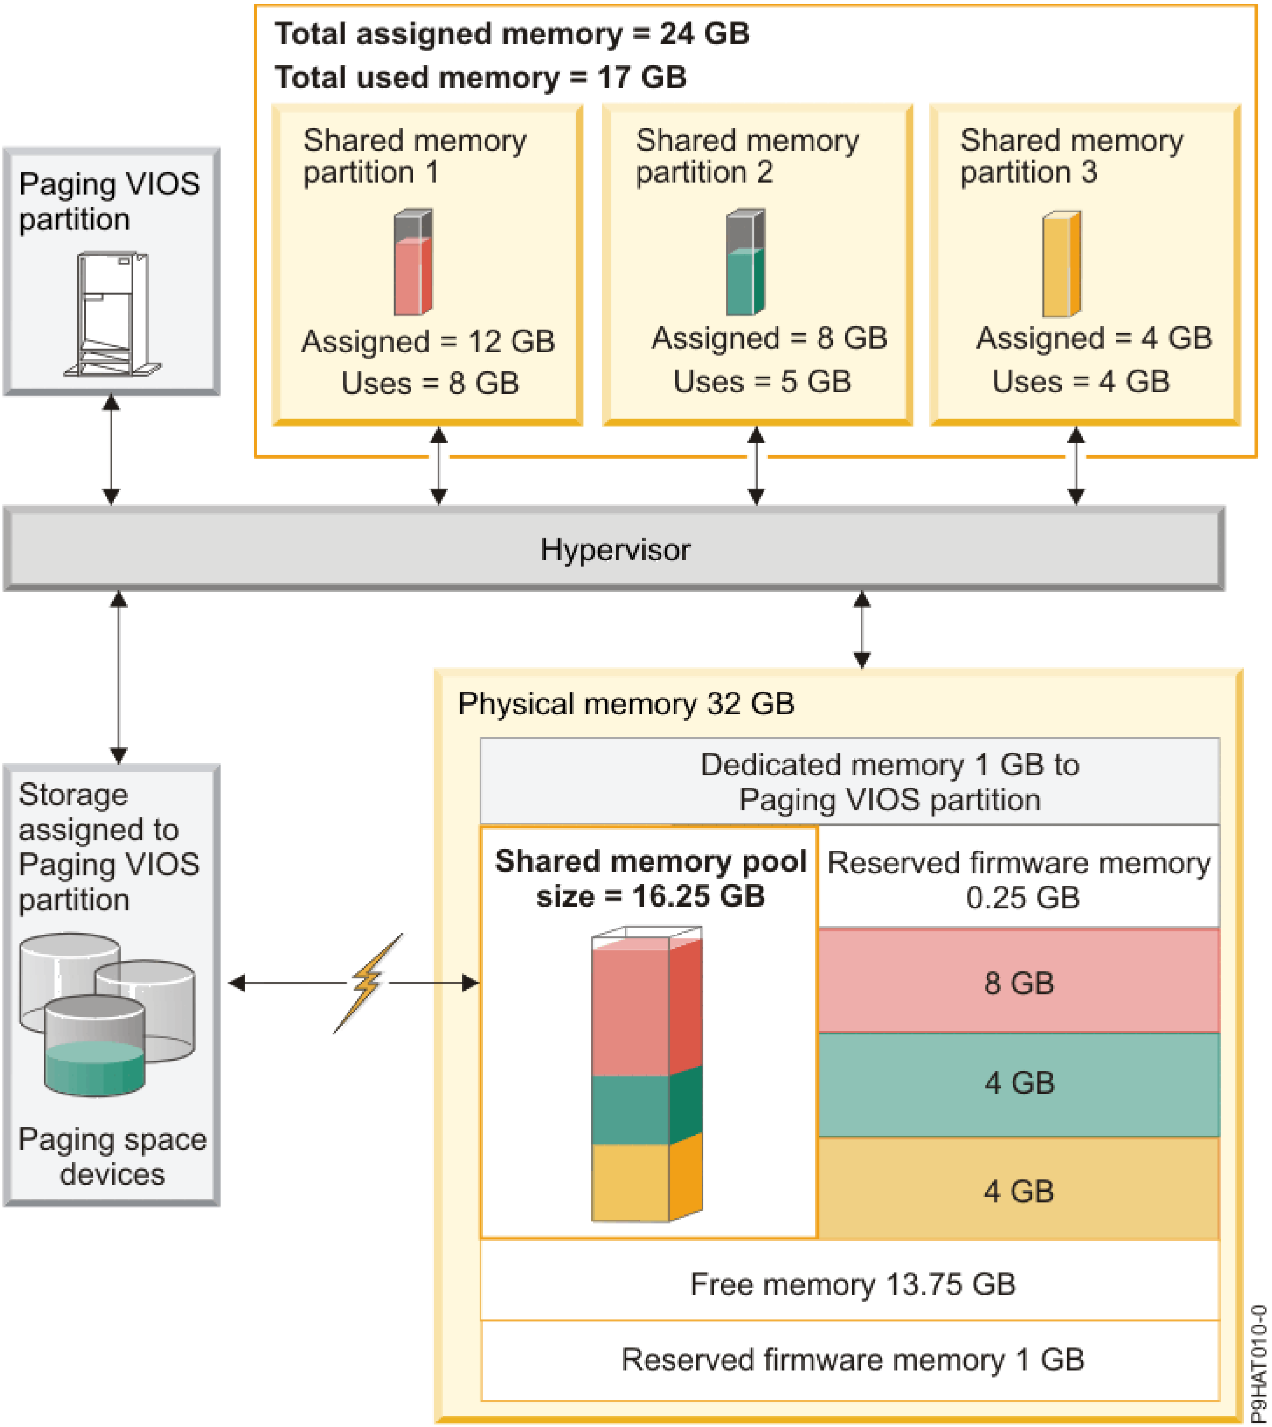
Total assigned memory = 24 GB
Total used memory = 17 GB
Paging VIOS partition
Shared memory partition 1
Assigned = 12 GB
Uses = 8 GB
Shared memory partition 2
Assigned = 8 GB
Uses = 5 GB
Shared memory partition 3
Assigned = 4 GB
Uses = 4 GB
Hypervisor
Physical memory 32 GB
Dedicated memory 1 GB to Paging VIOS partition
Storage assigned to Paging VIOS partition
Shared memory pool size = 16.25 GB
Reserved firmware memory 0.25 GB
8 GB
4 GB
4 GB
Paging space devices
Free memory 13.75 GB
Reserved firmware memory 1 GB
P9HAT010-0

Figure 2. A server with shared memory configuration that is physically over committed

The figure shows a shared memory pool of 16.25 GB that is shared among three shared memory partitions. The hypervisor uses a small portion (0.25 GB) of the shared memory pool to manage the shared memory resources. The figure also shows one paging VIOS partition that owns all of the physical storage in the system. The physical storage contains a paging space device for each shared memory partition. The paging VIOS partition does not use the memory in the shared memory pool, but rather receives dedicated memory of 1 GB. Of the remaining system memory, 1 GB is reserved for the hypervisor so that it can manage other system resources, and 13.75 GB is free memory that is available for system growth. For example, you can dynamically add more memory to the shared memory pool or you can create additional dedicated memory partitions.

Shared memory partition 1 is assigned 12 GB of logical memory, Shared memory partition 2 is assigned 8 GB of logical memory, and Shared memory partition 3 is assigned 4 GB of logical memory. Together, the shared memory partitions are assigned 24 GB of logical memory, which is more than the 16.25 GB of logical memory allocated to the shared memory pool. Therefore, the memory configuration is over committed.

Shared memory partition 1 currently uses 8 GB of physical memory, shared memory partition currently uses 5 GB of physical memory, and shared memory partition 3 currently uses 4 GB of physical memory. Together, the shared memory partitions currently use 17 GB of physical memory, which is greater than the amount of physical memory available to them in the shared memory pool. Therefore, the memory configuration is physically overcommitted. In other words, the shared memory pool does not contain enough physical memory for the hypervisor to satisfy the memory needs of all the shared memory partitions without storing some of the memory in the paging space devices. In this example, the difference of 1 GB is stored in the paging space device that is assigned to Shared memory partition 2. When Shared memory partition 2 needs to access data, the hypervisor might need to retrieve the data from the paging space device before the operating system can access it.

Related concepts

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Data flow for shared memory partitions

When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

In a shared memory configuration that is physically overcommitted (where the sum of the logical memory that is currently used by all the shared memory partitions is greater than the amount of memory in the shared memory pool), the hypervisor stores some of the logical memory that belongs to a shared memory partition in the shared memory pool and some of the logical memory in a paging space device. For the operating system in a shared memory partition to access its memory, the memory must be in the shared memory pool. Thus, when the operating system needs to access data that is stored on the paging space device, the hypervisor works with a paging VIOS partition to move the data from the paging space device to the shared memory pool so that the operating system can access it.

The following figure shows the data flow for shared memory.

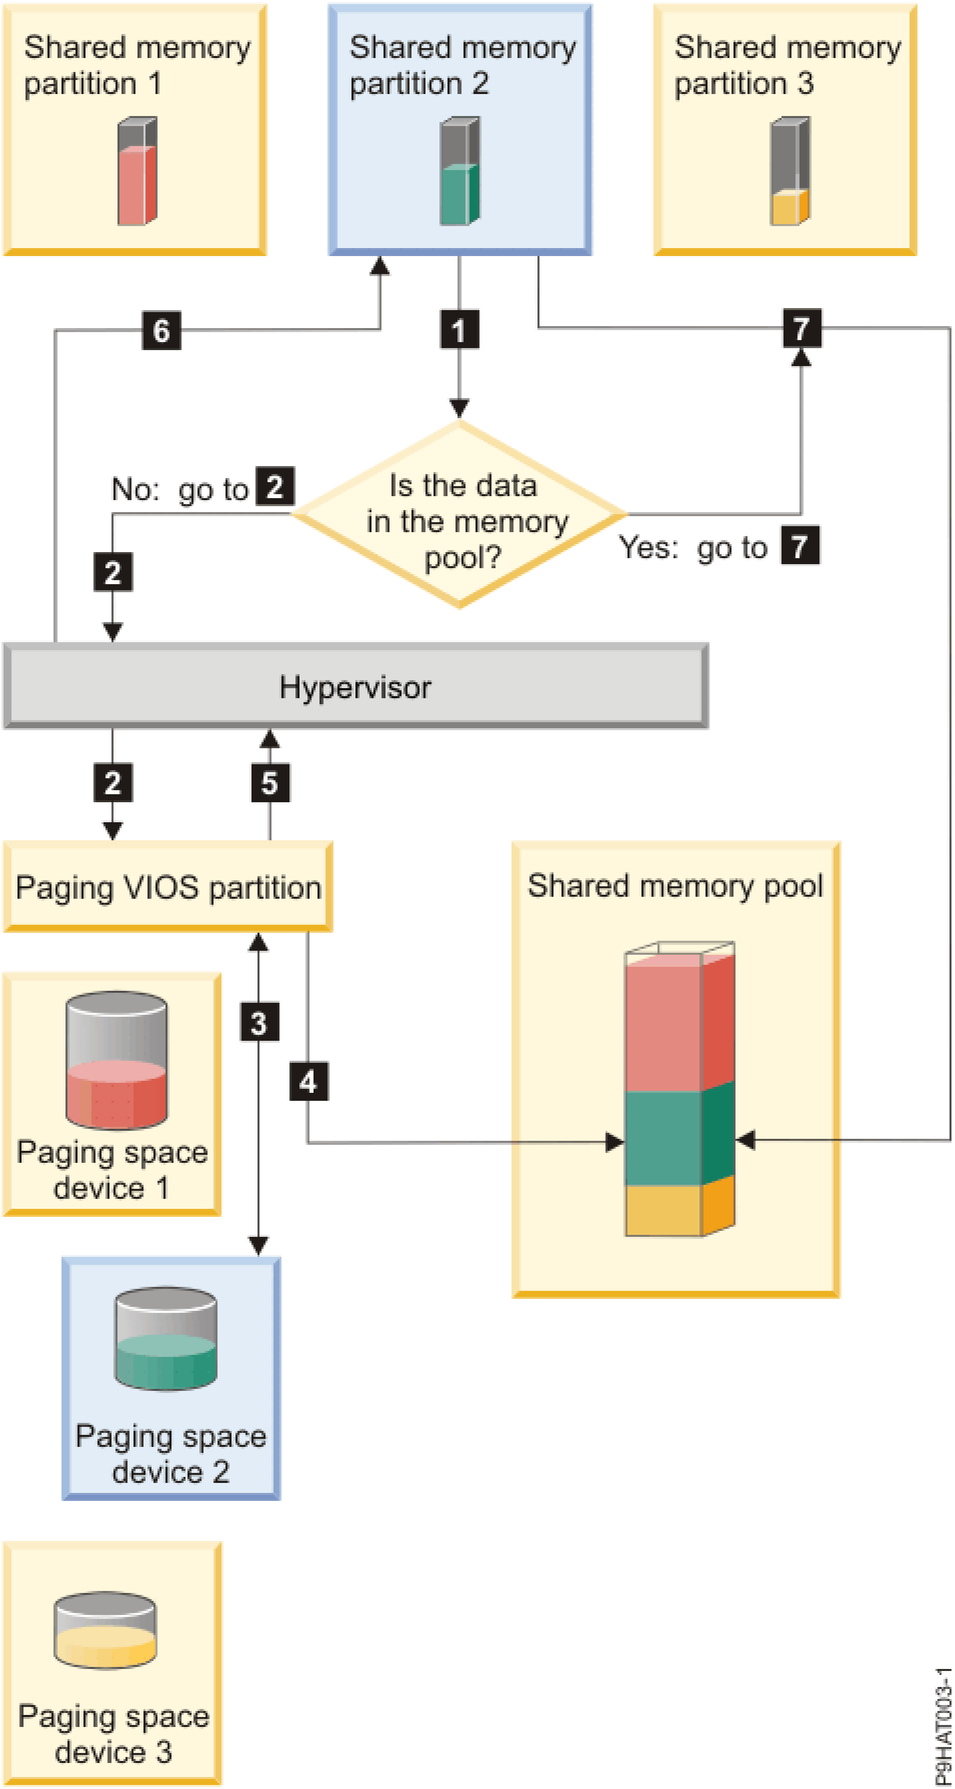

Figure 3. The process of managing data in a shared memory configuration that is overcommitted

In general, the data flows as follows:

1. The operating system that runs in a shared memory partition attempts to access data.
   - If the data is in the shared memory pool, processing continues with step "7" on page 28.
   - If the data is not in the shared memory pool, a page fault occurrs. The hypervisor inspects the page fault and discovers that the hypervisor moved the data to the paging space device, thereby causing the page fault. Processing continues with step "2" on page 28. (If the operating system that runs in the shared memory partition moved the data to auxiliary storage, thereby causing the page fault, then the operating system must retrieve the data.)
2. The hypervisor sends a request to a paging VIOS partition to retrieve the data from the paging space device and to write it to the shared memory pool.
3. The paging VIOS partition searches the paging space device that is assigned to the shared memory partition and finds the data.
4. The paging VIOS partition writes the data to the shared memory pool.
5. The paging VIOS partition notifies the hypervisor that the data is in the shared memory pool.
6. The hypervisor notifies the operating system that it can access the data.
7. The operating system accesses the data in the shared memory pool.

Related concepts

Logical memory
Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Paging space device
You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

Shared memory distribution
The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Logical memory

Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

You can configure minimum, maximum, desired, and assigned logical memory sizes for a shared memory partition.

Table 4. Logical memory sizes

Logical memory size	Description
Minimum	The minimum amount of logical memory with which you want the shared memory partition to operate. You can dynamically remove logical memory from the shared memory partition down to this value.
Maximum	The maximum amount of logical memory that the shared memory partition is allowed to use. You can dynamically add logical memory to the shared memory partition up to this value.

Table 4. Logical memory sizes (continued)

Logical memory size	Description
Desired	The amount of logical memory with which you want the shared memory partition to activate.
Assigned	The amount of logical memory that the shared memory partition can use. A shared memory partition does not have to use all of its assigned logical memory at any given time.

On systems that are managed by a Hardware Management Console (HMC), you configure the minimum, maximum, and desired logical memory sizes in the partition profile. When you activate the shared memory partition, the HMC assigns the desired logical memory to the shared memory partition.

The following figure shows a shared memory partition with its logical memory.

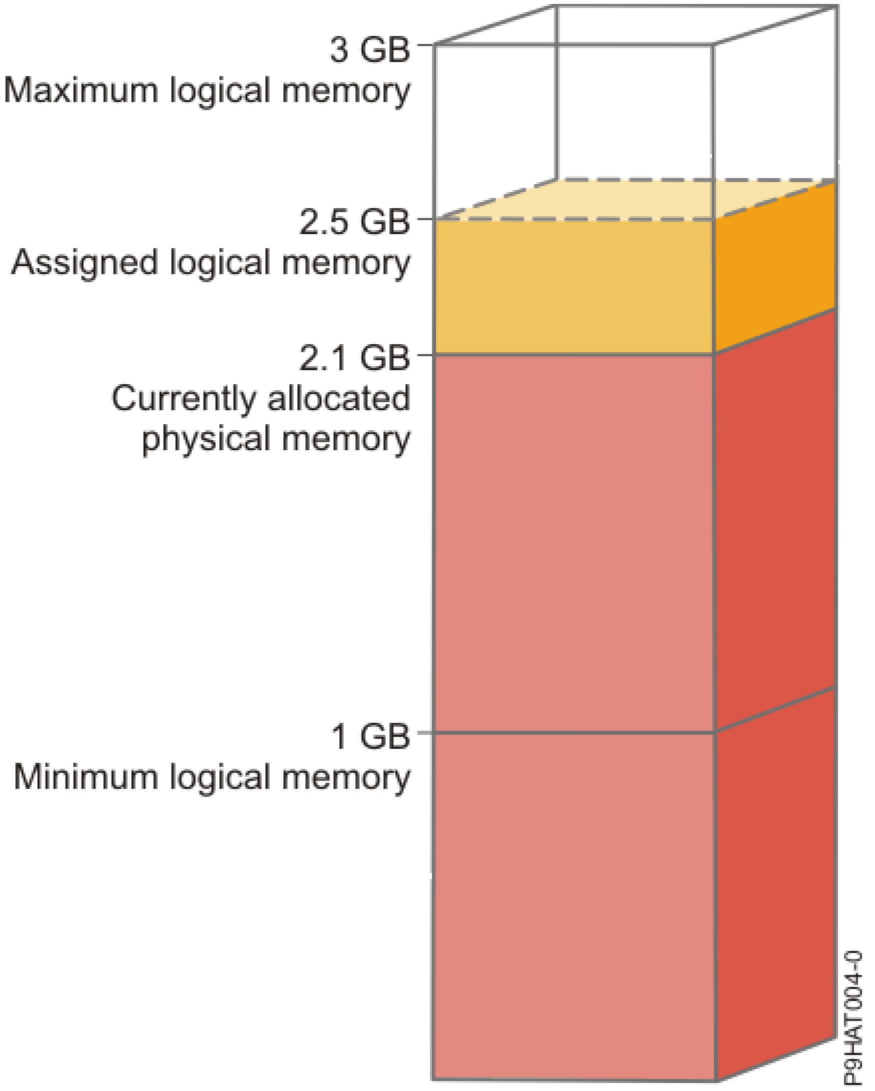

Figure 4. A shared memory partition that is assigned more logical memory than the amount of physical memory currently allocated to it

The figure shows a shared memory partition that is assigned 2.5 GB of logical memory. Its maximum logical memory is 3 GB and its minimum logical memory is 1 GB. You can change the assigned logical memory by dynamically adding or removing logical memory to or from the shared memory partition. You can dynamically add logical memory to the shared memory partition up to the maximum logical memory size, and you can dynamically remove logical memory from the shared memory partition down to its minimum logical memory size.

The figure also shows that the amount of physical memory that is currently allocated to the shared memory partition from the shared memory pool is 2.1 GB. If the workload that runs in the shared memory partition currently uses 2.1 GB of memory and requires an additional 0.2 GB of memory, and the shared memory pool is logically overcommitted, the hypervisor allocates an additional 0.2 GB of physical memory to the shared memory partition by assigning memory pages that are not currently in use by other shared memory partitions. If the shared memory pool is physically overcommitted, the hypervisor stores 0.2 GB of the shared memory partition's memory in a paging space device. When the shared memory partition needs to access the data that resides in the paging space device, the hypervisor retrieves the data for the operating system.

The amount of physical memory allocated to the shared memory partition can be less than the minimum logical memory size. This is because the minimum logical memory size is a boundary for logical memory, not for physical memory. In addition to the minimum logical memory size, the maximum, desired, and assigned logical memory sizes also do not control the amount of physical memory assigned to the shared memory partition. Likewise, dynamically adding or removing logical memory to or from a shared memory partition does not change the amount of physical memory allocated to the shared memory partition. When you set the logical memory sizes and dynamically add or remove logical memory, you set or change the amount of memory that the operating system can use, and the hypervisor decides how to distribute that memory between the shared memory pool and the paging space device.

Related concepts

Data flow for shared memory partitions
When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

Paging space device
You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

Shared memory distribution
The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Partition profile
A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Related tasks

Preparing to configure shared memory
Before you configure the shared memory pool and create logical partitions that use shared memory (hereafter referred to as *shared memory partitions*), you need to plan for the shared memory pool,

the shared memory partitions, the paging space devices, and the Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*).

I/O entitled memory
I/O entitled memory is the maximum amount of physical memory (from the shared memory pool) that is guaranteed to be available to a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) for its I/O devices at any given time.

Each shared memory partition is entitled to some portion of the shared memory pool so that the I/O devices that are assigned to the shared memory partition have access to physical memory during I/O operations. If the minimum amount of memory that I/O devices require for I/O operations does not reside in the shared memory pool for the duration the device needs the memory, the device fails. Virtual adapters that are entitled to physical memory from the shared memory pool include virtual SCSI adapters, virtual Ethernet adapters, and virtual Fibre Channel adapters. Virtual serial adapters are not entitled to physical memory from the shared memory pool.

The following figure shows a shared memory partition with I/O entitled memory.

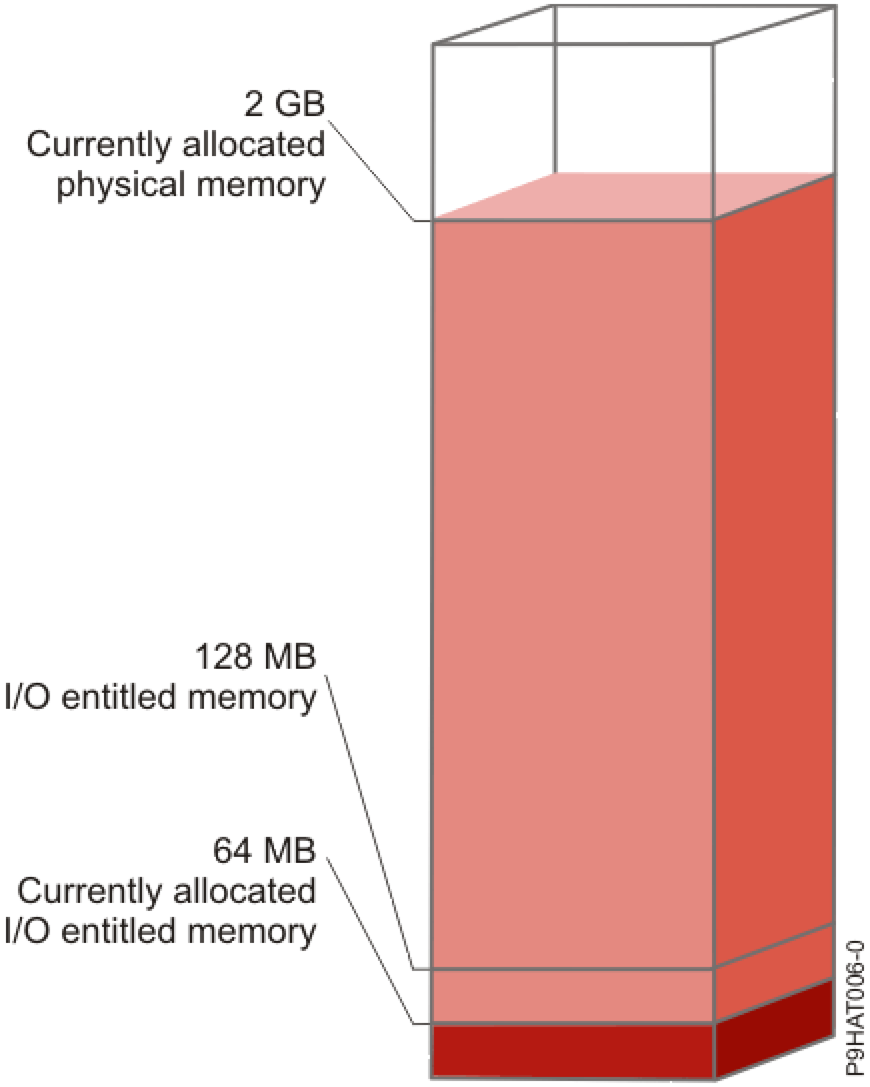

Figure 5. A shared memory partition whose I/O entitled memory is greater than the amount of physical memory that it currently uses for its I/O devices

The figure shows a shared memory partition with 128 MB of I/O entitled memory. The shared memory partition uses 64 MB of physical memory for its I/O devices, which is less than its I/O entitled memory of 128 MB.

As depicted in the previous figure, a shared memory partition might not use all of its I/O entitled memory at any given time. Unused portions of the I/O entitled memory assigned to a shared memory partition are available to the hypervisor to allocate to other shared memory partitions, if necessary. The hypervisor does not reserve unused portions of I/O entitled memory for the shared memory partition to use in the future. However, the hypervisor guarantees that the shared memory partition can use the entire portion of the I/O entitled memory that is assigned to it as needed. If the shared memory partition later requires some of its unused I/O entitled memory, the hypervisor must allocate enough physical memory from the shared memory pool to satisfy the new I/O memory requirement, without exceeding the I/O entitled memory that is assigned to the shared memory partition.

For example, you assign 128 MB of I/O entitled memory to a shared memory partition. The shared memory partition uses only 64 MB for its I/O devices. Thus, the hypervisor allocates 64 MB of physical memory from the shared memory pool to the shared memory partition for its I/O devices. The remaining 64 MB is available to the hypervisor to allocate to other shared memory partitions, if necessary. Later, you add two virtual adapters to the shared memory partition, each requiring 16 MB of memory. Thus, the shared memory partition needs an additional 32 MB of physical memory for its I/O devices. Because the shared memory partition currently uses only 64 MB of physical memory for its I/O devices and the shared memory partition is entitled to use up to 128 MB for its I/O devices, the hypervisor allocates an additional 32 MB of physical memory from the shared memory pool to the shared memory partition to accommodate the new virtual adapters. The shared memory partition now uses 96 MB of physical memory from the shared memory pool for its I/O devices.

Because unused portions of I/O entitled memory are available to the hypervisor to allocate elsewhere, for the amount of total physical memory that the hypervisor allocates from the shared memory pool to a shared memory partition can be less than the I/O entitled memory of the shared memory partition. The following figure shows this situation.

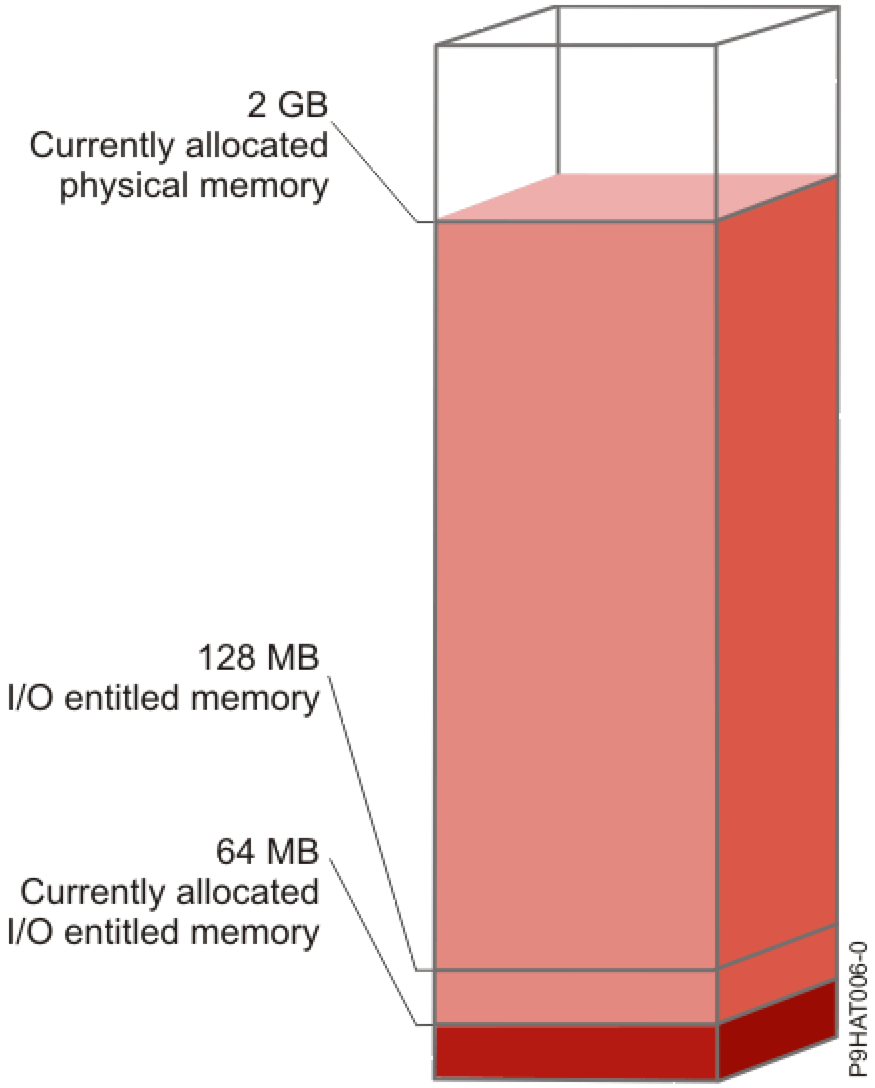

Figure 6. A shared memory partition whose I/O entitled memory is greater than the total amount of physical memory allocated to it

The figure shows a shared memory partition with 128 MB of I/O entitled memory. The shared memory partition uses 64 MB of physical memory for its I/O devices. The unused portion of the I/O entitled memory, 64 MB, is available to the hypervisor to allocate to other shared memory partitions, if necessary. The hypervisor allocates a total of 96 MB of physical memory from the shared memory pool to the shared memory partition, which is less than the I/O entitled memory of 128 MB.

When you create a shared memory partition, the Hardware Management Console (HMC) automatically sets the I/O entitled memory for the shared memory partition. When you activate a shared memory partition, the HMC sets the I/O entitled memory mode to the *auto* mode. In the auto mode, the HMC automatically adjusts the I/O entitled memory for the shared memory partition when you add or remove virtual adapters.

The I/O entitled memory mode can also be set to the *manual* mode. You can dynamically change the I/O entitled memory mode to the manual mode and then dynamically change the I/O entitled memory for the shared memory partition. When you add or remove a virtual adapter to or from the shared memory partition in manual mode, the HMC does not automatically adjust the I/O entitled memory. Therefore, you might need to dynamically adjust the I/O entitled memory when you dynamically add or remove adapters to or from the shared memory partition. On HMC-managed systems, you use the graphical interface to dynamically change the I/O entitled memory mode. When the I/O entitled memory mode is in the manual mode, you can also use the graphical interface to dynamically change the amount of I/O entitled memory that is assigned to a shared memory partition. When the I/O entitled memory mode is in the manual mode, you can also use the **chhwres** command to dynamically change the amount of I/O entitled memory that is assigned to a shared memory partition. When you restart a shared memory partition, the I/O entitled memory mode is set to the auto mode regardless of what the I/O entitled memory mode was set to before you restarted the shared memory partition.

When the amount of physical memory that a shared memory partition uses for its I/O devices is equal to the I/O entitled memory that is assigned to the shared memory partition, the shared memory partition cannot use any more physical memory for its I/O devices. In this situation, the following actions can occur:

- The operating system that runs in the shared memory partition manages the I/O operations so that the workload that runs in the shared memory partition operates within the I/O entitled memory that is assigned to the shared memory partition. If the workload attempts to use more physical memory for I/O operations than the I/O entitled memory that is assigned to the shared memory partition, the operating system delays some I/O operations while it runs other I/O operations. In this situation, the I/O entitled memory of the shared memory partition constrains the I/O configuration of the shared memory partition because the operating system does not have enough physical memory to run all of the I/O operations simultaneously.
- When you dynamically add a virtual adapter to the shared memory partition and the I/O entitled memory mode is in the manual mode, the I/O configuration of the shared memory partition might become constrained, or the adapter might fail when you attempt to configure it. If the adapter fails, enough I/O entitled memory is not assigned to the shared memory partition to accommodate the new adapter. To resolve the problem, you can dynamically increase the amount of I/O entitled memory that is assigned to the shared memory partition, or you can remove some existing virtual adapters from the shared memory partition. When you remove virtual adapters from the shared memory partition, the physical memory that those adapters were using becomes available for the new adapter.
- When you dynamically add a virtual adapter to the shared memory partition and the I/O entitled memory mode is in the auto mode, the HMC automatically increases the I/O entitled memory assigned to the shared memory partition to accommodate the new adapter. If the HMC cannot increase the I/O entitled memory of the shared memory partition, enough physical memory is not available in the shared memory pool for the hypervisor to allocate to the shared memory partition and the adapter cannot be assigned to the shared memory partition. To resolve the problem, you can add physical memory to the shared memory pool, or you can remove some existing virtual adapters from the shared memory partition. When you remove virtual adapters from the shared memory partition, the physical memory that those adapters were using becomes available for the new adapter.

Paging VIOS partition
A Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*) provides access to the paging space devices for the logical partitions that are assigned to the shared memory pool (hereafter referred to as *shared memory partitions*).

When the operating system that runs in a shared memory partition attempts to access data, and the data is located in the paging space device that is assigned to the shared memory partition, the hypervisor sends a request to a paging VIOS partition to retrieve the data and write it to the shared memory pool so that the operating system can access it.

A paging VIOS partition is not a shared memory partition and does not use the memory in the shared memory pool. A paging VIOS partition provides access to the paging space devices for the shared memory partitions.

HMC

On systems that are managed by a Hardware Management Console (HMC), you can assign one or two paging VIOS partitions to the shared memory pool. When you assign a single paging VIOS partition to the shared memory pool, the paging VIOS partition provides access to all of the paging space devices for the shared memory partitions. The paging space devices can be located in physical storage in the server or on a storage area network (SAN). When you assign two paging VIOS partitions to the shared memory pool, you can configure each paging VIOS partition to access paging space devices in one of the following ways:

- You can configure each paging VIOS partition to access independent paging space devices. Paging space devices that are accessed by only one paging VIOS partition, or independent paging space devices, can be located in physical storage in the server or on a SAN.
- You can configure both paging VIOS partitions to access the same, or common, paging space devices. In this configuration, the paging VIOS partitions provide redundant access to paging space devices. When one paging VIOS partition becomes unavailable, the hypervisor sends a request to the other paging VIOS partition to retrieve the data on the paging space device. Common paging space devices must be located on a SAN to enable symmetrical access from both paging VIOS partitions.
- You can configure each paging VIOS partition to access some independent paging space devices and some common paging space devices.

If you configure the shared memory pool with two paging VIOS partitions, you can configure a shared memory partition to use either a single paging VIOS partition or redundant paging VIOS partitions. When you configure a shared memory partition to use redundant paging VIOS partitions, you assign a primary paging VIOS partition and a secondary paging VIOS partition to the shared memory partition. The hypervisor uses the primary paging VIOS partition to access the shared memory partition's paging space device. At this point, the primary paging VIOS partition is the current paging VIOS partition for the shared memory partition. The current paging VIOS partition is the paging VIOS partition that the hypervisor uses at any point in time to access data in the paging space device that is assigned to the shared memory partition. If the primary paging VIOS partition becomes unavailable, the hypervisor uses the secondary paging VIOS partition to access the shared memory partition's paging space device. At this point, the secondary paging VIOS partition becomes the current paging VIOS partition for the shared memory partition and continues as the current paging VIOS partition even after the primary paging VIOS partition becomes available again.

You do not need to assign the same primary and secondary paging VIOS partitions to all of the shared memory partitions. For example, you assign paging VIOS partition A and paging VIOS partition B to the shared memory pool. For one shared memory partition, you can assign paging VIOS partition A as the primary paging VIOS partition and paging VIOS partition B as the secondary paging VIOS partition. For a different shared memory partition, you can assign paging VIOS partition B as the primary paging VIOS partition and paging VIOS partition A as the secondary paging VIOS partition.

The following figure shows an example of a system with four shared memory partitions, two paging VIOS partitions, and four paging space devices.

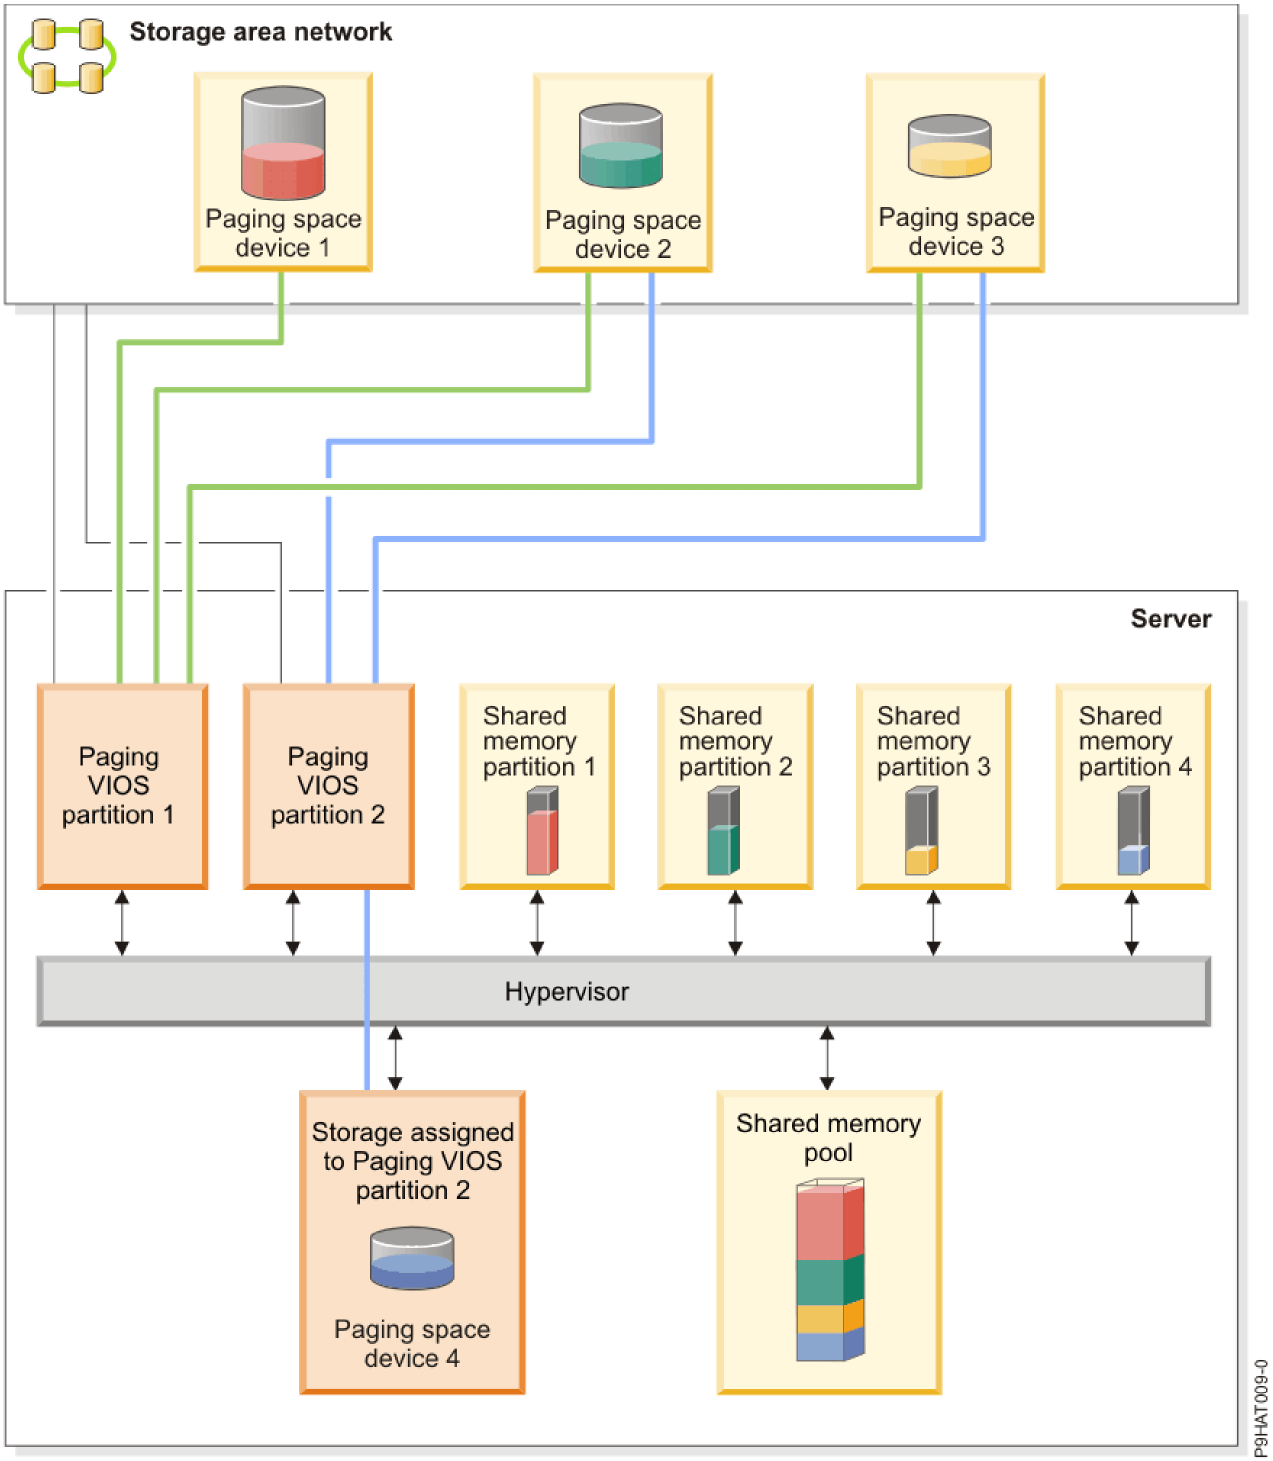

The example shows the configuration options for paging VIOS partitions and paging space devices as described in the following table.

Table 5. Examples of paging VIOS partition configurations

Configuration option	**Example**
The paging space device that is assigned to a shared memory partition is located in physical storage in the server and is accessed by a single paging VIOS partition.	Paging space device 4 provides the paging space for Shared memory partition 4. Shared memory partition 4 is assigned to use Paging VIOS partition 2 to access Paging space device 4. Paging space device 4 is located in physical storage in the server and is assigned to Paging VIOS partition 2. Paging VIOS partition 2 is the only paging VIOS partition that can access Paging space device 4 (This relationship is shown by the blue line that connects Paging VIOS partition 2 to Paging space device 4.).
The paging space device that is assigned to a shared memory partition is located on a SAN and is accessed by a single paging VIOS partition.	Paging space device 1 provides the paging space for Shared memory partition 1. Shared memory partition 1 is assigned to use Paging VIOS partition 1 to access Paging space device 1. Paging space device 1 is connected to the SAN. Paging VIOS partition 1 is also connected to the SAN and is the only paging VIOS partition that can access Paging space device 1 (This relationship is shown by the green line that connects Paging VIOS partition 1 to Paging space device 1.).

Table 5. Examples of paging VIOS partition configurations (continued)

Configuration option	Example
The paging space device that is assigned to a shared memory partition is located on a SAN and is accessed redundantly by two paging VIOS partitions.	Paging space device 2 provides the paging space for Shared memory partition 2. Paging space device 2 is connected to the SAN. Paging VIOS partition 1 and Paging VIOS partition 2 are also connected to the SAN and can both access Paging space device 2. (These relationships are shown by the green line that connects Paging VIOS partition 1 to Paging space device 2 and the blue line that connects Paging VIOS partition 2 to Paging space device 2.) Shared memory partition 2 is assigned to use redundant paging VIOS partitions to access Paging space device 2. Paging VIOS partition 1 is configured as the primary paging VIOS partition and Paging VIOS partition 2 is configured as the secondary paging VIOS partition. Similarly, Paging space device 3 provides the paging space for Shared memory partition 3. Paging space device 3 is connected to the SAN. Paging VIOS partition 1 and Paging VIOS partition 2 are also connected to the SAN and can both access Paging space device 3. (These relationships are shown by the green line that connects Paging VIOS partition 1 to Paging space device 3 and the blue line that connects Paging VIOS partition 2 to Paging space device 3.) Shared memory partition 3 is assigned to use redundant paging VIOS partitions to access Paging space device 3. Paging VIOS partition 2 is configured as the primary paging VIOS partition and Paging VIOS partition 1 is configured as the secondary paging VIOS partition. Because Paging VIOS partition 1 and Paging VIOS partition 2 both have access to Paging space device 2 and Paging space device 3, Paging space device 2 and Paging space device 3 are common paging space devices that are accessed redundantly by Paging VIOS partition 1 and Paging VIOS partition 2. If Paging VIOS partition 1 becomes unavailable and Shared memory partition 2 needs to access data on its paging space device, the hypervisor sends a request to Paging VIOS partition 2 to retrieve the data on Paging space device 2. Similarly, if Paging VIOS partition 2 becomes unavailable and Shared memory partition 3 needs to access the data on its paging space device, the hypervisor sends a request to Paging VIOS partition 1 to retrieve the data on Paging space device 3.

Table 5. Examples of paging VIOS partition configurations (continued)

Configuration option	Example
A paging VIOS partition accesses both independent and common paging space devices.	Paging space device 1 and Paging space device 4 are independent paging space devices because only one paging VIOS partition accesses each. Paging VIOS partition 1 accesses Paging space device 1, and Paging VIOS partition 2 accesses Paging space device 4. Paging space device 2 and paging space device 3 are common paging space devices because both paging VIOS partitions access each. (These relationships are shown by the green and blue lines that connect the paging VIOS partitions to the paging space devices.) Paging VIOS partition 1 accesses the independent paging space device Paging space device 1, and also accesses the common paging space devices Paging space device 2 and Paging space device 3. Paging VIOS partition 2 accesses the independent paging space device Paging space device 4 and also accesses the common paging space devices Paging space device 2 and Paging space device 3.

When a single paging VIOS partition is assigned to the shared memory pool, you must shut down the shared memory partitions before you shut down the paging VIOS partition so that the shared memory partitions are not suspended when they attempt to access their paging space devices. When two paging VIOS partitions are assigned to the shared memory pool and the shared memory partitions are configured to use redundant paging VIOS partitions, you do not need to shut down the shared memory partitions to shut down a paging VIOS partition. When one paging VIOS partition is shut down, the shared memory partitions use the other paging VIOS partition to access their paging space devices. For example, you can shut down a paging VIOS partition and install VIOS updates without shutting down the shared memory partitions.

You can configure multiple VIOS logical partitions to provide access to paging space devices. However, you can only assign up to two of those VIOS partitions to the shared memory pool at any given time.

After you configure the shared memory partitions, you can later change the redundancy configuration of the paging VIOS partitions for a shared memory partition by modifying the partition profile of the shared memory partition and restarting the shared memory partition with the modified partition profile:

- You can change which paging VIOS partitions are assigned to a shared memory partition as the primary and secondary paging VIOS partitions.
- You can change the number of paging VIOS partitions that are assigned to a shared memory partition.

Paging space device
You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

A *paging space device* is a physical or logical device that is used by a Virtual I/O Server to provide the paging space for a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*). The *paging space* is an area of nonvolatile storage used to hold portions of the shared memory partition's memory that are not resident in the shared memory pool.

Related concepts

Data flow for shared memory partitions
When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O

Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

Logical memory
Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Paging space devices on systems that are managed by an HMC
Learn about the location requirements, size requirements, and redundancy preferences for paging space devices on systems that are managed by a Hardware Management Console (HMC).

When you configure the shared memory pool, you assign can paging space devices to the shared memory pool. Paging space devices can be located in physical storage in the server or on a storage area network (SAN) as follows:

- Paging space devices that are accessed by a single Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) can be located in physical storage in the server or on a SAN.
- Paging space devices that are accessed redundantly by two paging VIOS partitions, or *common* paging space devices, must be located on a SAN.

When you activate a shared memory partition, the HMC allocates a paging space device (that is assigned to the shared memory pool) to the shared memory partition. The HMC allocates only one paging space device to a shared memory partition at a time. When you shut down a shared memory partition, its paging space device becomes available to the HMC to allocate elsewhere. Thus, the fewest number of paging space devices that must be assigned to the shared memory pool is equal to the number of shared memory partitions that you plan to run simultaneously. After you create the shared memory pool, you can add or remove paging space devices to or from the shared memory pool as needed.

The HMC assigns paging space devices to shared memory partitions based on the size requirements for the shared memory partition and the redundancy preferences that you specify for partition activation.

Size requirements

The HMC allocates a paging space device to a shared memory partition that best fits the size requirements of the shared memory partition.

- For AIX and Linux shared memory partitions, the paging space device must be at least the size of the maximum logical memory size of the shared memory partition.
- For IBM i shared memory partitions, the paging space device must be at least the size of the maximum logical memory size of the shared memory partition plus 8 KB for every megabyte.

Shared memory partitions might have several partition profiles that specify different maximum logical memory sizes. To maintain flexibility, consider creating paging space devices that are large enough to be used by shared memory partitions with multiple partition profiles. When you activate a shared memory partition with a different partition profile, the shared memory partition already has a paging space device allocated to it based on the size requirements of the previously activated partition profile. If you create a paging space device that is large enough to meet the size requirements of multiple partition profiles and you activate the shared memory partition with a different partition profile, the HMC can use the same paging space device for the newly activated partition profile. If the paging space device does not meet the size requirements of the newly activated partition profile, the HMC frees the paging space device currently allocated to the shared memory partition and allocates a different paging space device that meets the size requirements specified in the newly activated partition profile.

Redundancy preferences

The HMC allocates a paging space device to a shared memory partition that satisfies the redundancy preferences that you specify for partition activation:

- If you specify that the shared memory partition uses redundant paging VIOS partitions, the HMC uses the following process to select a suitable paging space device for the shared memory partition:
    1. The HMC assigns a paging space device that is common and available. (A paging space device is *available* when it is not currently assigned to a shared memory partition and is inactive.)
    2. If the HMC cannot find a paging space device that is common and available, it reassigns a paging space device that is common and unavailable. (A paging space device is *unavailable* when it is active and currently assigned to a shared memory partition that is shut down.)
    3. If the HMC cannot find a paging space device that is common and unavailable, it cannot activate the shared memory partition.
- If you specify that the shared memory partition does not use redundant paging VIOS partitions, the HMC uses the following process to select a suitable paging space device for the shared memory partition:
    1. The HMC assigns a paging space device that is independent and available. (A paging space device is *independent* when it is accessed by only one paging VIOS partition that is assigned to the shared memory partition.)
    2. If the HMC cannot find a paging space device that is independent and available, the HMC reassigns a paging space device that is independent and unavailable.
    3. If the HMC cannot find a paging space device that is independent and unavailable, and two paging VIOS partitions are assigned to the shared memory pool, then the HMC assigns a paging space device that is common and available. In this situation, the shared memory partition does not use redundant paging VIOS partitions even though its paging space device can be accessed by both paging VIOS partitions. Also, the partition profile need not specify the second paging VIOS partition.
    4. If the HMC cannot find a paging space device that is common and available, and two paging VIOS partitions are assigned to the shared memory pool, then the HMC reassigns a paging space device that is common and unavailable. In this situation, the shared memory partition does not use redundant paging VIOS partitions even though its paging space device can be accessed by both paging VIOS partitions. Also, the partition profile need not specify the second paging VIOS partition.
    5. If the HMC cannot find a paging space device that is common and unavailable, it cannot activate the shared memory partition.
- If you specify that the shared memory partition use redundant paging VIOS partitions, if possible, the HMC uses the following process to select a suitable paging space device for the shared memory partition:
    1. The HMC assigns a paging space device that is common and available.
    2. If the HMC cannot find a paging space device that is common and available, it assigns a paging space device that is common and unavailable.
    3. If the HMC cannot find a paging space device that is common and unavailable, it assigns a paging space device that is independent and available to the primary paging VIOS partition. In this situation, the shared memory partition does not use redundant paging VIOS partitions and the primary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.
    4. If the HMC cannot find a paging space device that is independent and available to the primary paging VIOS partition, it assigns a paging space device that is independent and unavailable to the primary paging VIOS partition. In this situation, the shared memory partition does not use redundant paging VIOS partitions and the primary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.
    5. If the HMC cannot find a paging space device that is independent and unavailable to the primary paging VIOS partition, it assigns a paging space device that is independent and available to the secondary paging VIOS partition. In this situation, the shared memory partition does not use redundant paging VIOS partitions and the secondary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.
    6. If the HMC cannot find a paging space device that is independent and available to the secondary paging VIOS partition, it assigns a paging space device that is independent and unavailable to the secondary paging VIOS partition. In this situation, the shared memory partition does not use

redundant paging VIOS partitions and the secondary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.

7. If the HMC cannot find a paging space device that is independent and unavailable to the secondary paging VIOS partition, it cannot activate the shared memory partition.

Related concepts

Partition profile
A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Related tasks

Preparing to configure shared memory
Before you configure the shared memory pool and create logical partitions that use shared memory (hereafter referred to as *shared memory partitions*), you need to plan for the shared memory pool, the shared memory partitions, the paging space devices, and the Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*).

Related reference

Configuration requirements for shared memory
Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

Shared memory distribution

The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

In a shared memory configuration that is physically overcommitted (where the sum of the logical memory that is currently used by of all shared memory partitions is greater than the amount of memory in the shared memory pool), the hypervisor stores a portion of the logical memory in the shared memory pool and stores the remainder of the logical memory in the paging space devices. The hypervisor determines the amount of physical memory to allocate from the shared memory pool to each shared memory partition and the amount of logical memory to store in the paging space devices. The hypervisor also determines which pieces, or pages, of memory to store in each location.

The smallest amount of physical memory that the hypervisor can allocate from the shared memory pool to a shared memory partition at any given time is the amount of physical memory that the shared memory partition requires for its I/O devices. The hypervisor guarantees to each shared memory partition that the shared memory partition can use a portion of the shared memory pool for its I/O devices, up to the I/O entitled memory that is assigned to the shared memory partition. The largest amount of physical memory that the hypervisor can allocate from the shared memory pool to a shared memory partition at any given time is the amount of logical memory assigned to the shared memory partition.

The amount of physical memory from the shared memory pool that the hypervisor allocates to the shared memory partitions is determined by the workloads that are running in the shared memory partitions and the amount of logical memory that is assigned to each shared memory partition. You can influence how much physical memory the hypervisor allocates from the shared memory pool to each shared memory partition by specifying a memory weight for each shared memory partition. *Memory weight* is a relative value that is one of the factors that the hypervisor uses to allocate physical memory from the shared memory pool to the shared memory partitions. A higher memory weight relative to the memory weights of other shared memory partitions increases the probability that the hypervisor allocates more physical memory to a shared memory partition.

To help maintain the best possible performance, the operating system that runs in a shared memory partition continually attempts to operate within the amount of physical memory allocated to it from the shared memory pool by moving its overcommitted logical memory to a paging space. In general, the operating system moves its memory to a paging space more often when it runs in a shared memory partition than when it runs in a dedicated memory partition. Therefore, the paging space that the operating system uses to manage its memory needs to be larger when the logical partition uses shared memory than when the logical partition uses dedicated memory.

The operating systems that run in shared memory partitions provide information to the hypervisor about how the operating system uses its pages. When the hypervisor manages the overcommitted logical memory, it uses this information to determine which pages to store in the paging space device and which pages to store in the shared memory pool. When the hypervisor needs to deallocate physical memory from the shared memory partition and move it to the paging space device, the hypervisor requests the operating system to release pages. The operating system might mark the pages that it will not use, and the hypervisor moves the marked pages first. This enables the hypervisor to select the most optimal pages to move out of the shared memory pool, which improves memory use and performance. For example, the operating system uses one page for kernel data and another page for cache and the hypervisor needs to move one page to the paging space device. The hypervisor moves the cache page to the paging space device to optimize performance.

Related concepts

Paging space device
You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

Data flow for shared memory partitions
When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Logical memory
Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Active Memory Expansion for AIX logical partitions

When you enable Active Memory Expansion for an AIX logical partition, you increase the memory capacity of the logical partition without assigning more memory to it. The operating system compresses a portion of the memory that the logical partition uses. This compression creates space for more data and expanding the memory capacity of the logical partition.

When you expand the memory capacity of a logical partition, you enable the logical partition to do more work with the same amount of memory. This can be especially useful when you want to increase the workload of a logical partition, but cannot assign more memory to the logical partition. When you expand the memory capacity of several logical partitions on a server, you increase the overall memory capacity of the server. This can be especially useful when you want to consolidate more workloads onto the server by creating more logical partitions.

You can configure the degree of memory expansion that you want to achieve for a logical partition by setting the Active Memory Expansion factor in a partition profile of the logical partition. The expansion factor is a multiplier of the amount of memory that is assigned to the logical partition. For example, if the amount of memory that is assigned to a logical partition is 25 GB and the expansion factor is set to 2.0, then the desired memory capacity of the logical partition is 50 GB. In this situation, the operating system attempts to compress data such that 50 GB of data fits into 25 GB of memory. After you set the expansion factor, you can monitor the performance of the logical partition and then dynamically change the expansion factor to improve performance.

When you configure a logical partition to use Active Memory Expansion, you must also configure some additional processing resources for the logical partition. The operating system uses the additional processing resources to perform the memory compression. The amount of processing resources that the logical partition requires depends on the workload that is running in the logical partition and the expansion factor that you set for the logical partition.

You can configure Active Memory Expansion for logical partitions that use dedicated memory and logical partitions that use shared memory.

Related information

IBM AIX Knowledge Center website

Terminal and console options for logical partitions

You can initiate a terminal or console session to the logical partitions on your managed system using various methods. Your choice of terminal or console depends on your operating system and business needs.

The following choices of terminal or console are available for each operating system.

Table 6. Terminal and console options for logical partitions

Operating system	Terminal or console options
AIX	• Hardware Management Console (HMC) • Telnet • OpenSSH with OpenSSL (included in the AIX expansion pack) • Direct serial connection (ASCII terminal or PC connected with null modem cable) • IBM i virtual console (for AIX logical partitions that use IBM i resources) • When on a system with a Virtual I/O Server (VIOS) logical partition, the console can be provided by the VIOS logical partition when using VIOS 1.2.0 or later.
Linux	• HMC • Telnet • OpenSSH with OpenSSL (included in the Linux distribution) • Direct serial connection (ASCII terminal or PC connected with null modem cable) • IBM i virtual console (for Linux logical partitions that use IBM i resources) • When on a system with a Virtual I/O Server (VIOS) logical partition, the console can be provided by the VIOS logical partition when using VIOS 1.2.0 or later.
Virtual I/O Server	• Hardware Management Console (HMC) • Telnet • OpenSSH with OpenSSL (included in the AIX expansion pack) • Direct serial connection (ASCII terminal or PC connected with null modem cable) • IBM i virtual console (for AIX logical partitions that use IBM i resources) • When on a system with a Virtual I/O Server (VIOS) logical partition, the console can be provided by the VIOS logical partition when using VIOS 1.2.0 or later.

Hardware Management Console terminal and console options

The HMC provides virtual terminal emulation for AIX and Linux logical partitions and virtual 5250 console emulation for IBM i logical partitions.

You can create virtual terminal and virtual 5250 console sessions locally on the HMC by using the Server Management commands on the HMC. If you configure the HMC to allow remote access, you can also create virtual terminal and virtual 5250 console sessions remotely through the HMC. You can create remote virtual terminal sessions on AIX and Linux logical partitions by using the Server Management commands. You can also create virtual 5250 console sessions on IBM i logical partitions. You must configure the HMC to allow remote access, and you must configure encryption on the logical partitions for the session to be secure.

The HMC communicates with servers by using service applications to detect, consolidate, and send information to IBM for analysis.

The following figure shows a partitioned server being managed by an HMC.

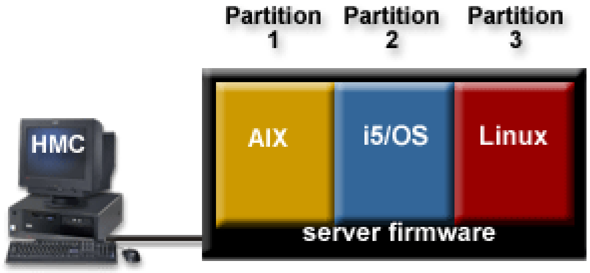

Operations Console for IBM i logical partitions

Operations Console allows you to use a local or remote PC to access IBM i on logical partitions. Operations Console is an installable component of IBM i Access for Windows licensed program.

You can use Operations Console to connect to IBM i logical partitions over a network (LAN connection).

The management tasks that you can perform using Operations Console depend upon whether you are managing your logical partitions using a Hardware Management Console (HMC) or using the Virtual Partition Manager on IBM i:

- If you are using an HMC to manage your logical partitions, you can use Operations Console to access IBM i on your IBM i logical partitions.
- If you are using the Virtual Partition Manager on IBM i to manage your logical partitions, you can use Operations Console to access IBM i on your IBM i logical partitions. In turn, you can use Operations Console to access the Virtual Partition Manager on your IBM i logical partitions. This allows you to create up to four Linux logical partitions on the managed system and manage the resources for all logical partitions on the managed system.

The following figure shows a partitioned server with an HMC and local console on a network.

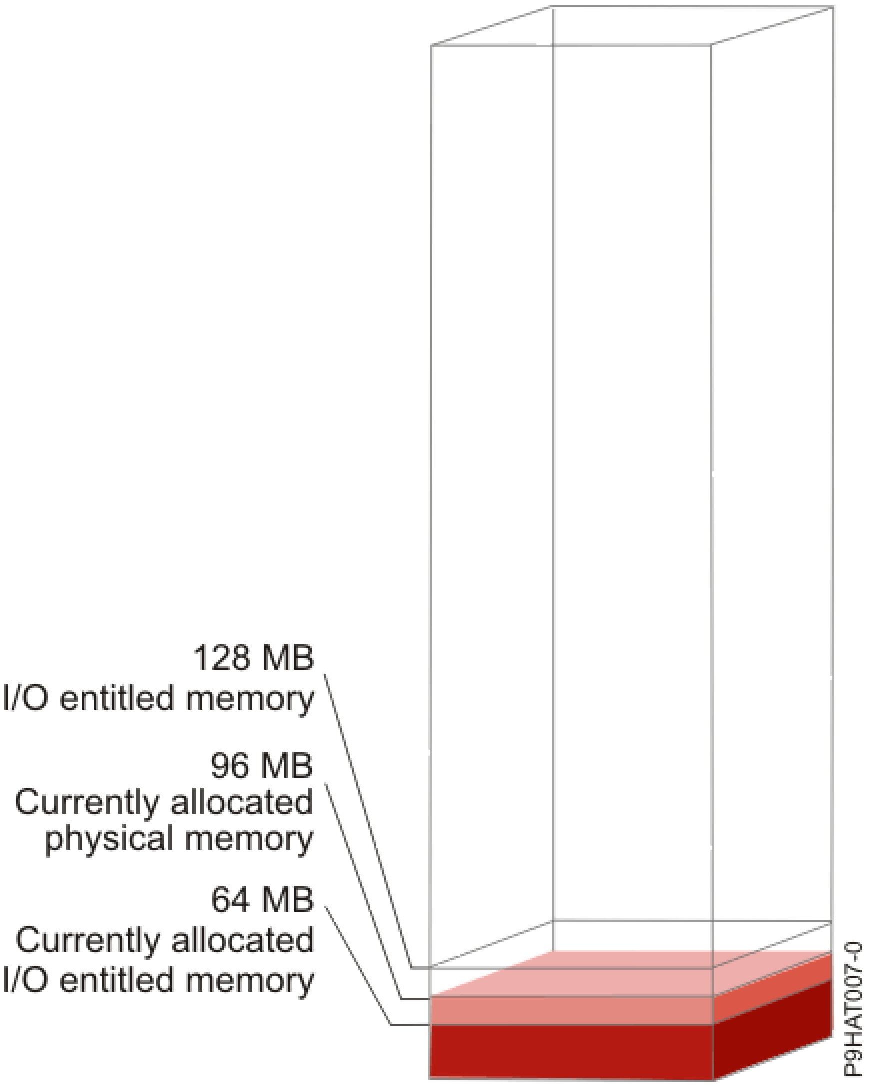

Related information

Managing Operations

I/O devices

I/O devices allow your managed system to gather, store, and transmit data. I/O devices are found in the server unit itself and in expansion units that are attached to the server. I/O devices can be embedded into the unit, or they can be installed into physical slots.

Not all types of I/O devices are supported for all operating systems or on all server models. For example, Switch Network Interface (SNI) adapters are supported only on certain server models, and are not supported for IBM i logical partitions.

Single root I/O virtualization (SR-IOV) allows virtualization of the physical ports of an adapter so that the ports can be shared by multiple partitions that are running simultaneously. To share the ports of an SR-IOV capable adapter, the adapter must first be enabled for the SR-IOV shared mode. After an adapter is enabled for SR-IOV shared mode, SR-IOV logical ports can be assigned to logical partitions.

⚠ **Attention:** Some PCI adapters and embedded controllers require multiple PCI or PCI-E slots to be associated with them. Carefully review the PCI or PCI-E slot assignments for each logical partition to ensure that the slot configuration of the logical partition meets the adapter functional requirements. For details, see Managing PCI adapters, and PCI adapter placement.

Virtual adapters

With virtual adapters, you can connect logical partitions with each other without using physical hardware. Operating systems can display, configure, and use virtual adapters just like they can display, configure, and use physical adapters. Depending on the operating environment used by the logical partition, you can create virtual Ethernet adapters, virtual Fibre Channel adapters, virtual Small Computer Serial Interface (SCSI) adapters, and virtual serial adapters for a logical partition.

The system administrator uses the following tools to create virtual adapters:

- Hardware Management Console (HMC)
- Virtual Partition Manager

Adapters can be added while the system is running using dynamic partitioning. The virtual adapters are recorded in system inventory and management utilities. Converged location codes can be used to correlate operating-system level or partition-level software entities to adapters, such as eth0, CMN21, and en0. Similarly, the Ethernet adapters are visible in the same way as physical Ethernet adapters.

By default, virtual Ethernet Media Access Control (MAC) addresses are created from the locally administered range. Using the default MAC addresses, it is possible that different servers will have virtual Ethernet adapters with the same addresses. This situation can present a problem if multiple, virtual networks are bridged to the same physical network.

If a server logical partition providing I/O for a client logical partition fails, the client logical partition might continue to function, depending on the significance of the hardware it is using. For example, if one logical partition is providing the paging volume for another logical partition, a failure of the logical partition providing that particular resource will be significant to the other logical partition. However, if the shared resource is a tape drive, a failure of the server logical partition providing the resource will have only minimal effects on the client logical partition.

Client support for virtual I/O

The following table summarizes operating system support for using virtual I/O devices.

Table 7. Client support for virtual I/O by operating system

Client operating system	Virtual console	Virtual Ethernet	Virtual Fibre Channel	Virtual disk	Virtual optical	Virtual tape
AIX	Yes	Yes	Yes	Yes	Yes, on HMC-managed systems, at least one Virtual I/O Server logical partition must be present	Yes
AIX	Yes	Yes	Yes	Yes	Yes, on HMC-managed systems, at least one Virtual I/O Server logical partition must be present	Yes
IBM i	Yes	Yes	Yes	Yes	Yes	Yes
Linux	Yes	Yes	Yes	Yes	Yes	Yes

Table 7. Client support for virtual I/O by operating system (continued)

Client operating system	Virtual console	Virtual Ethernet	Virtual Fibre Channel	Virtual disk	Virtual optical	Virtual tape
Linux	Yes	Yes	Yes	Yes	Yes	Yes

AIX logical partitions support booting from virtual devices, including disk boot from virtual disk, network boot from virtual Ethernet, and tape boot from virtual tape.

The firmware running in AIX and Linux logical partitions recognizes virtual I/O and can start the logical partition from virtual I/O. An IPL can be performed either from the network over virtual Ethernet or from a device such as virtual disk or virtual CD.

Server support for virtual I/O

The following table summarizes operating system support for providing virtual I/O to logical partitions.

Table 8. Server support for virtual I/O by operating system

Server	Virtual optical	Virtual console	Virtual disk	Virtual tape	Virtual Fibre Channel
IBM i	Yes	Yes	Yes	Yes	No
Linux	Yes	Yes	No	No	No
Linux	Yes	Yes	No	No	No
Virtual I/O Server	Yes	Yes	Yes	Yes	Yes

Virtual I/O Server provides SCSI disk, shared Ethernet, virtual Fibre Channel, virtual optical, and virtual tape functions to logical partitions that use Virtual I/O Server resources. On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, you can create a cluster of only one Virtual I/O Server (VIOS) partition connected to the same shared storage pool and that has access to distributed storage. On VIOS Version 2.2.1.3, or later, you can create a cluster that consists of up to four VIOS partitions. The Virtual I/O Server also provides a virtual console to AIX and Linux logical partitions.

IBM i provides disk, CD, tape, and console functions to logical partitions that use IBM i resources. IBM i uses standard network server storage and network server descriptions to provide disk, CD, and tape resources to other logical partitions. An IBM i logical partition cannot simultaneously provide virtual resources to other logical partitions and use virtual resources provided by another IBM i logical partition or by the Virtual I/O Server logical partition.

To configure virtual I/O for the logical partitions on your managed system, you must create virtual I/O adapters on the HMC. Virtual I/O adapters are usually created when you create your logical partitions. Alternatively, you can add virtual I/O adapters to running logical partitions using dynamic partitioning. After you create a virtual I/O adapter, you can then access the operating system used by the logical partition and complete the configuration of the virtual I/O adapter in the operating system software. For Linux partitions, virtual adapters are listed in the device tree. The device tree contains virtual SCSI adapters, not the devices under the adapter.

Logical Host Ethernet Adapter

A logical Host Ethernet Adapter (LHEA) is a special type of virtual adapter. Even though an LHEA is a virtual resource, an LHEA can exist only if a physical Host Ethernet Adapter, or Integrated Virtual Ethernet, provides its resources to the LHEA.

Note: HEA is not supported on POWER9 processor-based server.

Related concepts

I/O device assignment in partition profiles
I/O devices are assigned to partition profiles either on a slot-by-slot basis, or on logical port basis in the case of shared mode single root I/O virtualization (SR-IOV) adapters. For I/O devices that are assigned to partition profiles on a slot-by-slot basis, most I/O devices can be assigned to a partition profile on the HMC as required or as allocated. For SR-IOV logical ports, I/O devices are always assigned to a profile as required.

Host Ethernet Adapter
A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Virtual Ethernet
Virtual Ethernet allows logical partitions to communicate with each other without having to assign physical hardware to the logical partitions.

You can create virtual Ethernet adapters on each logical partition and connect these virtual Ethernet adapters to virtual LANs. TCP/IP communications over these virtual LANs is routed through the server firmware.

A virtual Ethernet adapter provides similar function as a 1 Gb Ethernet adapter. A logical partition can use virtual Ethernet adapters to establish multiple high-speed interpartition connections within a single managed system. AIX, IBM i, Linux, and Virtual I/O Server logical partitions can communicate with each other using TCP/IP over the virtual Ethernet communications ports.

Virtual Ethernet adapters are connected to an IEEE 802.1q (VLAN)-style virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VLAN IDs that enable them to share a common logical network. The virtual Ethernet adapters are created and the VLAN ID assignments are done using the Hardware Management Console (HMC). The system transmits packets by copying the packet directly from the memory of the sender logical partition to the receive buffers of the receiver logical partition without any intermediate buffering of the packet.

You can configure an Ethernet bridge between the virtual LAN and a physical Ethernet adapter that is owned by a Virtual I/O Server or IBM i logical partition. The logical partitions on the virtual LAN can communicate with an external Ethernet network through the Ethernet bridge. You can reduce the number of physical Ethernet adapters required for a managed system by routing external communications through the Ethernet bridge.

The number of virtual Ethernet adapters allowed for each logical partition varies by operating system.

- AIX 5.3 and later supports up to 256 virtual Ethernet adapters for each logical partition.
- Version 2.6 of the Linux kernel supports up to 32, 768 virtual Ethernet adapters for each logical partition. Each Linux logical partition can belong to a maximum of 4, 094 virtual LANs.

Besides a Port VLAN ID, the number of additional VLAN ID values that can be assigned for each virtual Ethernet adapter is 19, which indicates that each virtual Ethernet adapter can be used to access 20 networks. The HMC generates a locally administered Ethernet MAC address for the virtual Ethernet adapters so that these addresses do not conflict with physical Ethernet adapter MAC addresses.

After a specific virtual Ethernet is enabled for a logical partition, a network device is created in the logical partition. This network device is named `ent`*X* on AIX logical partitions and `eth`*X* on Linux logical partitions, where *X* represents sequentially assigned numbers. The user can then set up TCP/IP configuration similar to a physical Ethernet device to communicate with other logical partitions.

If a virtual Ethernet adapter is set for checksum offload, the virtual Ethernet adapter cannot generate a checksum for any packet that the virtual Ethernet adapter sends to a multicast or broadcast MAC address.

Some managed systems contain a Host Ethernet Adapter (HEA). A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters). Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA using an Ethernet bridge on another logical partition.

Note: HEA is not supported on POWER9 processor-based server.

You can enable and disable individual virtual Ethernet adapters by using the Hardware Management Console (HMC). You can use the **chhwres** command to enable or disable a virtual Ethernet adapter. A particular logical partition can be removed from the network when the virtual Ethernet adapter is disabled. You can reconnect the logical partition to the network by enabling the virtual Ethernet adapter. To reconnect the logical partition, you must use a virtual Ethernet that is bridged by using a Shared Ethernet Adapter (SEA) in the Virtual I/O Server (VIOS). The status of the virtual Ethernet adapter can be queried at any time by using the **lshwres** command. The disabled state persists during partition restart. Trunk adapters cannot be disabled. You must have super administrator or product engineer access to the HMC to enable or disable a virtual Ethernet adapter.

Related concepts

Host Ethernet Adapter
A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Virtual Fibre Channel
With N_Port ID Virtualization (NPIV), you can configure the managed system so that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter.

To access physical storage in a typical storage area network (SAN) that uses Fibre Channel, the physical storage is mapped to logical units (LUNs) and the LUNs are mapped to the ports of physical Fibre Channel adapters. Each physical port on each physical Fibre Channel adapter is identified using one worldwide port name (WWPN).

NPIV is a standard technology for Fibre Channel networks that enables you to connect multiple logical partitions to one physical port of a physical Fibre Channel adapter. Each logical partition is identified by a unique WWPN, which means that you can connect each logical partition to independent physical storage on a SAN.

To enable NPIV on the managed system, you must create a Virtual I/O Server logical partition (version 2.1, or later) that provides virtual resources to client logical partitions. You assign the physical Fibre Channel adapters (that support NPIV) to the Virtual I/O Server logical partition. Then, you can connect virtual Fibre Channel adapters on the client logical partitions to virtual Fibre Channel adapters on the Virtual I/O Server logical partition. A *virtual Fibre Channel adapter* is a virtual adapter that provides client logical partitions with a Fibre Channel connection to a storage area network through the Virtual I/O Server logical partition. The Virtual I/O Server logical partition provides the connection between the virtual Fibre Channel adapters on the Virtual I/O Server logical partition and the physical Fibre Channel adapters on the managed system.

The following figure shows a managed system configured to use NPIV.

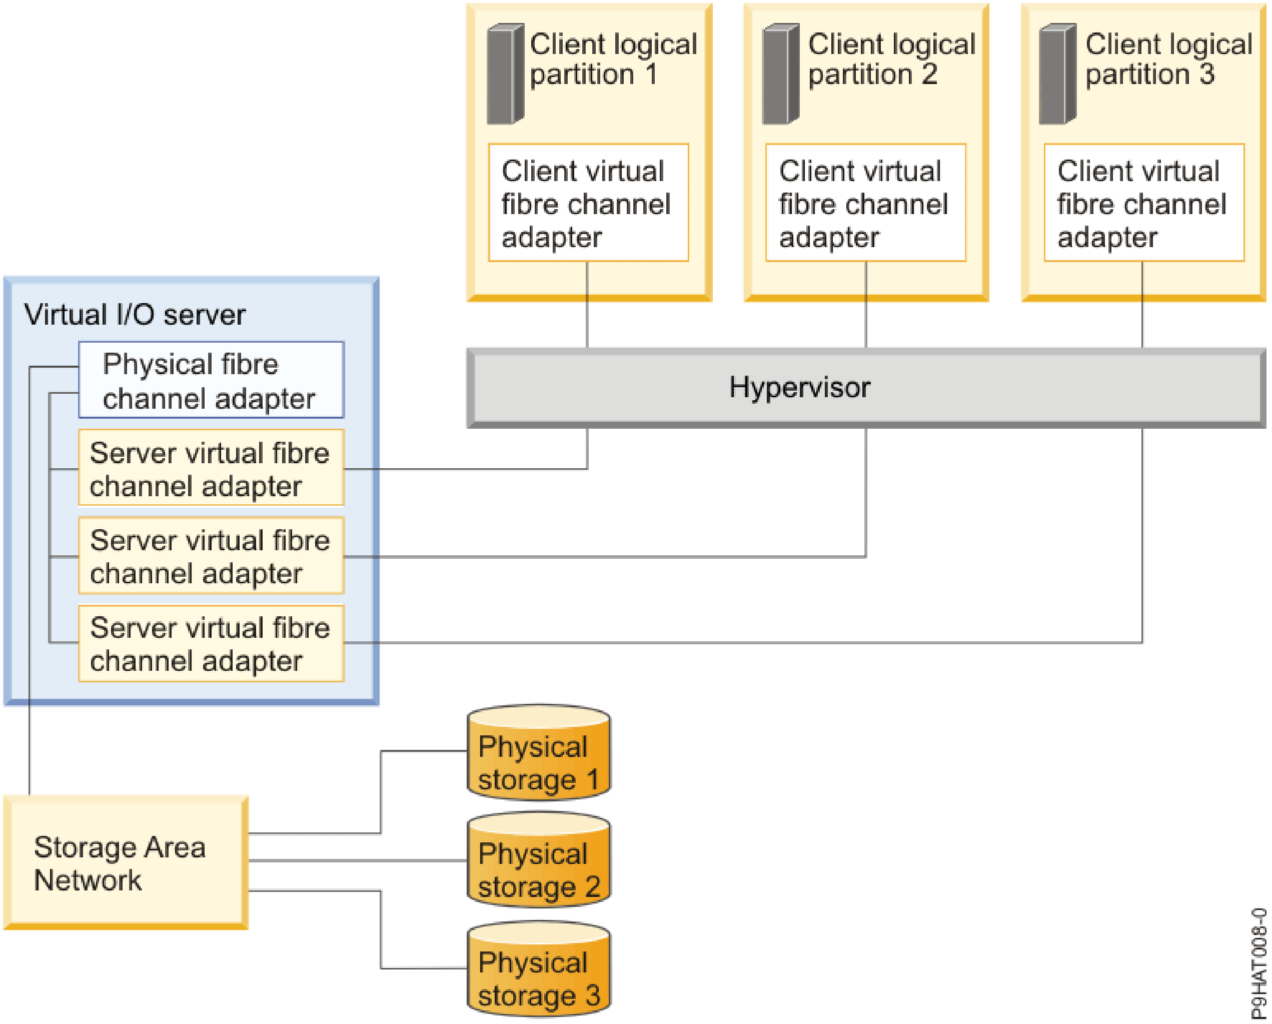
Client logical partition 1
Client virtual fibre channel adapter
Client logical partition 2
Client virtual fibre channel adapter
Client logical partition 3
Client virtual fibre channel adapter
Virtual I/O server
Physical fibre channel adapter
Server virtual fibre channel adapter
Server virtual fibre channel adapter
Server virtual fibre channel adapter
Hypervisor
Physical storage 1
Physical storage 2
Physical storage 3
Storage Area Network
P9HAT008-0

The figure shows the following connections:

- A storage area network (SAN) connects three units of physical storage to a physical Fibre Channel adapter that is located on the managed system. The physical Fibre Channel adapter is assigned to the Virtual I/O Server and supports NPIV.
- The physical Fibre Channel adapter connects to three virtual Fibre Channel adapters on the Virtual I/O Server. All three virtual Fibre Channel adapters on the Virtual I/O Server connect to the same physical port on the physical Fibre Channel adapter.
- Each virtual Fibre Channel adapter on the Virtual I/O Server connects to one virtual Fibre Channel adapter on a client logical partition. Each virtual Fibre Channel adapter on each client logical partition receives a pair of unique WWPNs. The client logical partition uses one WWPN to log into the SAN at any given time. The other WWPN is used when you move the client logical partition to another managed system.

Using their unique WWPNs and the virtual Fibre Channel connections to the physical Fibre Channel adapter, the operating systems that run in the client logical partitions discover, instantiate, and manage their physical storage located on the SAN. In the previous figure, Client logical partition 1 accesses Physical storage 1, Client logical partition 2 accesses Physical storage 2, and Client logical partition 3 accesses Physical storage 3. For IBM i client partitions, the LUNs of the physical storage connected with NPIV require a storage-specific device driver and do not use the generic virtual SCSI device driver. The Virtual I/O Server cannot access and does not emulate the physical storage to which the client logical partitions have access. The Virtual I/O Server provides the client logical partitions with a connection to the physical Fibre Channel adapters on the managed system.

There is always a one-to-one relationship between virtual Fibre Channel adapters on the client logical partitions and the virtual Fibre Channel adapters on the Virtual I/O Server logical partition. That is, each virtual Fibre Channel adapter on a client logical partition must connect to only one virtual Fibre Channel adapter on the Virtual I/O Server logical partition, and each virtual Fibre Channel on the Virtual I/O Server logical partition must connect to only one virtual Fibre Channel adapter on a client logical partition. Mapping of multiple Virtual Fibre Channel adapters of a single client logical partition through multiple virtual server Fibre Channel adapters to the same physical Fibre Channel adapter is not recommended.

Using SAN tools, you can zone and mask LUNs that include WWPNs that are assigned to virtual Fibre Channel adapters on client logical partitions. The SAN uses WWPNs that are assigned to virtual Fibre Channel adapters on client logical partitions the same way it uses WWPNs that are assigned to physical ports.

Related information

Redundancy configuration using Virtual Fibre channel adapters

Virtual Fibre Channel for HMC-managed systems

On systems that are managed by the Hardware Management Console (HMC), you can dynamically add and remove virtual Fibre Channel adapters to and from the Virtual I/O Server logical partition and each client logical partition. You can also view information about the virtual and physical Fibre Channel adapters and the worldwide port names (WWPNs) by using Virtual I/O Server commands.

To enable N_Port ID Virtualization (NPIV) on the managed system, you create the required virtual Fibre Channel adapters and connections as follows:

- You use the HMC to create virtual Fibre Channel adapters on the Virtual I/O Server logical partition and associate them with virtual Fibre Channel adapters on the client logical partitions.
- You use the HMC to create virtual Fibre Channel adapters on each client logical partition and associate them with virtual Fibre Channel adapters on the Virtual I/O Server logical partition. When you create a virtual Fibre Channel adapter on a client logical partition, the HMC generates a pair of unique WWPNs for the client virtual Fibre Channel adapter.
- You can connect the virtual Fibre Channel adapters on the Virtual I/O Server to the physical ports of the physical Fibre Channel adapter by running the **vfcmap** command on the Virtual I/O Server.

The HMC generates WWPNs based on the range of names available for use with the prefix in the vital product data on the managed system. This 6–digit prefix comes with the purchase of the managed system and includes 32,000 pairs of WWPNs. When you remove a virtual Fibre Channel adapter from a client logical partition, the hypervisor deletes the WWPNs that are assigned to the virtual Fibre Channel adapter on the client logical partition. The HMC does not reuse the deleted WWPNs when generating WWPNs for virtual Fibre Channel adapters in the future. If you run out of WWPNs, you must obtain an activation code that includes another prefix with another 32,000 pairs of WWPNs.

To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client logical partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client logical partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.

You can dynamically add and remove virtual Fibre Channel adapters to and from the Virtual I/O Server logical partition and to and from client logical partitions.

Table 9. Dynamic partitioning tasks and results for virtual Fibre Channel adapters

Dynamically add or remove virtual Fibre Channel adapter	To or from a client logical partition or a Virtual I/O Server logical partition	Result
Add a virtual Fibre Channel adapter	To a client logical partition	The HMC generates the a pair of unique WWPNs for the client virtual Fibre Channel adapter.
Add a virtual Fibre Channel adapter	To a Virtual I/O Server logical partition	You need to connect the virtual Fibre Channel adapter to a physical port on a physical Fibre Channel adapter.

Table 9. Dynamic partitioning tasks and results for virtual Fibre Channel adapters (continued)

Dynamically add or remove virtual Fibre Channel adapter	To or from a client logical partition or a Virtual I/O Server logical partition	Result
Remove a virtual Fibre Channel adapter	From a client logical partition	• The hypervisor deletes the WWPNs and does not reuse them. • You must either remove the associated virtual Fibre Channel adapter from the Virtual I/O Server, or associate it with another virtual Fibre Channel adapter on a client logical partition.
Remove a virtual Fibre Channel adapter	From a Virtual I/O Server logical partition	• The Virtual I/O Server removes the connection to the physical port on the physical Fibre Channel adapter. • You must either remove the associated virtual Fibre Channel adapter from the client logical partition, or associate it with another virtual Fibre Channel adapter on the Virtual I/O Server logical partition.

The following table lists the Virtual I/O Server commands that you can run to view information about the Fibre Channel adapters.

Table 10. Virtual I/O Server commands that display information about Fibre Channel adapters

Virtual I/O Server command	Information displayed by command
lsmap	• Displays the virtual Fibre Channel adapters on the Virtual I/O Server that are connected to the physical Fibre Channel adapter • Displays attributes of the virtual Fibre Channel adapters on the client logical partitions that are associated with the virtual Fibre Channel adapters on the Virtual I/O Server that are connected to the physical Fibre Channel adapter
lsnports	Displays information about the physical ports on the physical Fibre Channel adapters that support NPIV, such as: • The name and location code of the physical port • The number of available physical ports • The total number of WWPNs that the physical port can support • Whether the switches, to which the physical Fibre Channel adapters are cabled, support NPIV

You can also run the **lshwres** command on the HMC to display the remaining number of WWPNs and to display the prefix that is currently used to generate the WWPNs.

Virtual SCSI adapters
Virtual SCSI (Small Computer Systems Interface) adapters provide one logical partition with the ability to use storage I/O (disk, CD, and tape) that is owned by another logical partition.

A virtual SCSI client adapter in one logical partition can communicate with a virtual SCSI server adapter in another logical partition. The virtual SCSI client adapter allows a logical partition to access a storage device being made available by the other logical partition. The logical partition owning the hardware is the *server logical partition*, and the logical partition that uses the virtualized hardware is the *client logical partition*. With this arrangement, the system can have many server logical partitions.

For example, logical partition A provides disk space to logical partitions B, C, and D. A logical partition can simultaneously use virtual I/O from more than one logical partition. Therefore, using the example, while logical partition A provides disk space to logical partitions B, C, and D, logical partitions A and B can use a tape drive connected to logical partition D. In this case, A is serving D for disk space, while D is serving A for the tape device.

Virtual SCSI allows you to simplify the backup and maintenance operations on your managed system. When you back up the data on the server logical partition, you also back up the data on each client logical partition.

Virtual SCSI server adapters can be created only in logical partitions of type IBM i and Virtual I/O Server.

The virtual SCSI client device driver is not capable of storage protection using Redundant Arrays of Independent Disks (RAID). While the Linux operating system allows software RAID protection of virtual disks, the recommended technique for protecting disk storage is to configure the virtual I/O storage server to perform the disk protection.

For HMC-managed systems, virtual SCSI adapters are created and assigned to logical partitions using partition profiles.

Related concepts

Partition profile
A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Virtual serial adapters
Virtual serial adapters provide a point-to-point connection from one logical partition to another, or from the Hardware Management Console (HMC) to each logical partition on the managed system. Virtual serial adapters are used primarily to establish terminal or console connections to logical partitions.

When you create a logical partition, the HMC automatically creates two virtual server serial adapters on the logical partition. These virtual server serial adapters allow you to establish a terminal or console connection to the logical partition through the HMC.

You can also create pairs of virtual serial adapters on logical partitions so that you can access and control one logical partition directly from another logical partition. For example, one logical partition uses the disk resources of another logical partition using virtual SCSI adapters. You can create a server serial adapter on the logical partition that uses the disk resources and a client serial adapter on the logical partition that owns the disk resources. This connection allows the logical partition that owns the disk resources to shut down the logical partition that uses the disk resources before you back up data on the logical partition that owns the disk resources.

On HMC-managed systems, virtual serial adapters are created and assigned to logical partitions using partition profiles.

Related concepts

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Host Ethernet Adapter

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Note: HEA is not supported on POWER9 processor-based server.

Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA without having to go through an Ethernet bridge on another logical partition.

To connect a logical partition to an HEA, you must create a logical Host Ethernet Adapter (LHEA) for the logical partition. A *logical Host Ethernet Adapter (LHEA)* is a representation of a physical HEA on a logical partition. An LHEA appears to the operating system as if it were a physical Ethernet adapter, just as a virtual Ethernet adapter appears as if it were a physical Ethernet adapter. When you create an LHEA for a logical partition, you specify the resources that the logical partition can use on the actual physical HEA. Each logical partition can have one LHEA for each physical HEA on the managed system. Each LHEA can have one or more logical ports, and each logical port can connect to a physical port on the HEA.

You can create an LHEA for a logical partition using either of the following methods:

- You can add the LHEA to a partition profile, shut down the logical partition, and reactivate the logical partition using the partition profile with the LHEA.
- You can add the LHEA to a running logical partition using dynamic partitioning. This method can be used for Linux logical partitions only if you install the following operating systems on the logical partition:
  - Red Hat® Enterprise Linux version 4.6, or later
  - Red Hat Enterprise Linux version 5.1, or later
  - SUSE Linux Enterprise Server Version 10, or later
  - SUSE Linux Enterprise Server Version 11, or later

When you activate a logical partition, the LHEAs in the partition profile are considered to be required resources. If the physical HEA resources required by the LHEAs are not available, then the logical partition cannot be activated. However, when the logical partition is active, you can remove any LHEAs you want from the logical partition. For every active LHEA that you assign to an IBM i logical partition, IBM i requires 40 MB of memory.

After you create an LHEA for a logical partition, a network device is created in the logical partition. This network device is named `ent`*X* on AIX logical partitions, `CMN`*XX* on IBM i logical partitions, and `eth`*X* on Linux logical partitions, where *X* represents sequentially assigned numbers. The user can then set up TCP/IP configuration like a physical Ethernet device to communicate with other logical partitions.

You can configure a logical partition so that it is the only logical partition that can access a physical port of an HEA by specifying *promiscuous mode* for an LHEA that is assigned to the logical partition. When an LHEA is in promiscuous mode, no other logical partitions can access the logical ports of the physical port that is associated with the LHEA that is in promiscuous mode. You might want to configure a logical partition to promiscuous mode in the following situations:

- If you want to connect more than 16 logical partitions to each other and to an external network through a physical port on an HEA, you can create a logical port on a Virtual I/O Server and configure an Ethernet bridge between the logical port and a virtual Ethernet adapter on a virtual LAN. This allows all logical partitions with virtual Ethernet adapters on the virtual LAN to communicate with the physical port through the Ethernet bridge. If you configure an Ethernet bridge between a logical port and a virtual Ethernet adapter, the physical port that is connected to the logical port must have the following properties:

  - The physical port must be configured so that the Virtual I/O Server is the promiscuous mode logical partition for the physical port.
  - The physical port can have only one logical port.
- You want the logical partition to have dedicated access to a physical port.
- You want to use tools such as `tcpdump` or `iptrace`.

A logical port can communicate with all other logical ports that are connected to the same physical port on the HEA. The physical port and its associated logical ports form a logical Ethernet network. Broadcast and multicast packets are distributed on this logical network as though it was a physical Ethernet network. You can connect up to 16 logical ports to a physical port using this logical network. By extension, you can connect up to 16 logical partitions to each other and to an external network through this logical network. The actual number of logical ports that you can connect to a physical port depends upon the Multi-Core Scaling value of the physical port group. It also depends on the number of logical ports that have been created for other physical ports within the physical port group. By default, the Multi-Core Scaling value of each physical port group is set to 4, which allows four logical ports to be connected to the physical ports in the physical port group. To allow up to 16 logical ports to be connected to the physical ports in the physical port group, you must change the Multi-Core Scaling value of the physical port group to 1 and restart the managed system.

You can set each logical port to restrict or allow packets that are tagged for specific VLANs. You can set a logical port to accept packets with any VLAN ID, or you can set a logical port to accept only the VLAN IDs that you specify. You can specify up to 20 individual VLAN IDs for each logical port.

The physical ports on an HEA are always configured on the managed system level. If you use an HMC to manage a system, you must use the HMC to configure the physical ports on any HEAs belonging to the managed system. Also, the physical port configuration applies to all logical partitions that use the physical port. (Some properties might require setup in the operating system as well. For example, the maximum packet size for a physical port on the HEA must be set on the managed system level using the HMC. However, you must also set the maximum packet size for each logical port within the operating system.) By contrast, if a system is unpartitioned and is not managed by an HMC, you can configure the physical ports on an HEA within the operating system as if the physical ports were ports on a regular physical Ethernet adapter.

HEA hardware does not support half-duplex mode.

You can change the properties of a logical port on an LHEA by using dynamic partitioning to remove the logical port from the logical partition. You can also add the logical port back to the logical partition using the changed properties. If the operating system of the logical partition does not support dynamic partitioning for LHEAs, and you want to change any logical port property other than the VLANs on which the logical port participates, you must set a partition profile for the logical partition so that the partition profile contains the wanted logical port properties, shut down the logical partition, and activate the logical partition using the new or changed partition profile. If the operating system of the logical partition does not support dynamic partitioning for LHEAs, and you want to change the VLANs on which the logical port participates, you must remove the logical port from a partition profile belonging to the logical partition, shut down and activate the logical partition using the changed partition profile, add the logical port back to the partition profile using the changed VLAN configuration, and shut down and activate the logical partition again using the changed partition profile.

Related concepts

Virtual Ethernet
Virtual Ethernet allows logical partitions to communicate with each other without having to assign physical hardware to the logical partitions.

Related information

Shared Ethernet Adapters

Integrated Virtual Ethernet Adapter Technical Overview and Introduction

Tagged resources for IBM i logical partitions

When you create an IBM i logical partition using the Hardware Management Console (HMC), you must tag I/O adapters (IOAs) to perform specific functions for the IBM i logical partition.

A *tagged resource* is an IOA that is selected because it controls a device that performs a specific function for a logical partition. The HMC and the IBM i operating system use this tagging to locate and use the correct I/O device for each I/O function. For example, when you create an IBM i partition profile, you must tag the I/O device that you want the IBM i logical partition to use as its load source. The tag allows the HMC to locate the load source when you activate the logical partition using the partition profile.

You can tag the IOA that controls the I/O device that you want to use. Tagging the IOA allows you to specify the exact I/O device that you want to use.

The following table lists and describes the device types that are tagged and indicates whether you must tag the device type for IBM i logical partitions.

Table 11. Devices associated with tagged IOAs

Device	Description	Tagging required for IBM i logical partitions?
Alternate restart device	This device can be a tape drive or an optical device. The media in the alternate restart device is what the system uses to start from when you perform a D-mode initial program load (IPL). The alternate restart device loads the Licensed Internal Code that is contained on the removable media instead of the code on the load source disk unit.	Yes
Logical partition console	The first workstation that the system activates in the logical partition and the only device it activates on a manual IPL. The logical partition assumes that a console is always available for use.	Yes (if you are using a console device other than the HMC)
Load source disk unit	Each IBM i logical partition must have 1 disk unit designated as the load source. The system uses the load source to start the logical partition. The system always identifies this disk unit as unit number 1.	Yes

If you use the Virtual Partition Manager to create logical partition on your managed system, you do not have to tag I/O devices for these I/O functions. The IBM i logical partition automatically owns all physical I/O resources on the managed system, and the Virtual Partition Manager automatically tags the I/O device to use for each I/O function. The Virtual Partition Manager tags I/O devices for I/O functions based upon server model and location within the server. If you are partitioning a new server using the Virtual Partition Manager and have ordered the server with preinstalled IBM i, then you do not have to verify the placement of I/O devices within your new server.

Related concepts

I/O device assignment in partition profiles

I/O devices are assigned to partition profiles either on a slot-by-slot basis, or on logical port basis in the case of shared mode single root I/O virtualization (SR-IOV) adapters. For I/O devices that are assigned to partition profiles on a slot-by-slot basis, most I/O devices can be assigned to a partition profile on the HMC as required or as allocated. For SR-IOV logical ports, I/O devices are always assigned to a profile as required.

Load source placement rules for IBM i logical partitions
You must properly place a disk unit within a system unit or expansion unit before you use the disk unit as the load source for an IBM i logical partition. The placement rules depend upon the server unit or expansion unit in which the load source is located and, sometimes on, the I/O adapter (IOA) that controls the load source.

Note: The information provided does not replace the System Planning Tool (SPT). Use this information as a resource with the SPT output. Its purpose is to assist you in the load source placement for your IBM i logical partitions.

- There is no specific slot requirement for serial-attached SCSI (SAS) disk units. Any slot can contain the load source.
- The load-source IOA must be specified when you create your logical partition.
- Disk compression must be disabled for the load source disk.
- Disk units must have at least 17 GB of usable capacity.
- Disk mirroring requires two load source disk devices in valid load-source positions.
- Any disk IOA that can attach to a system capable of having logical partitions can be used for additional storage capacity after the special requirements for the load-source disk are met.
- Each logical partition has its own single-level storage and hence its own ASP configuration. The ASP configuration rules that apply to systems without logical partitions also apply to logical partitions.
- Disk protection can be defined for a logical partition in the same way as for a nonpartitioned system: parity protection (RAID), mirroring, or mixed. Bus-level mirroring requires two buses in the logical partition.
- Disk units that are already in use by a logical partition cannot be easily added to a different logical partition. You must first remove them from the configuration of the logical partition that is using the disk units before you add them to a different logical partition. In doing so, the system automatically moves any user or system data to other disk units in the same ASP.

Alternate restart device placement rules for IBM i logical partitions
You can use the internal optical device in the system unit, or you can use an external tape or optical device, to load the Licensed Internal Code and IBM i to the load source disk unit of an IBM i logical partition.

The only supported internal device for alternate restart is the Slimline DVD drive in the removable media slot in the system unit.

IBM i logical partitions have the following rules for the external alternate restart devices:

- The alternate restart device must be connected to bus 0 or port 0 of the IOA.
- The alternate restart IOA is specified during logical partition setup.

Switchable devices for IBM i logical partitions

When you set up an I/O adapter (IOA) so that it can be switched from one logical partition to another, you can share the devices that are associated with that IOA among many IBM i logical partitions.

When you switch an IOA, you take the control of the devices away from one logical partition and give it to another without restarting the server or the logical partition. Before switching the IOA to another logical partition, you must ensure that the device is not in use.

IOAs that are good candidates for switching between logical partitions include IOAs that are attached to high-cost devices or low-use or low-demand devices.

⚠ **Attention:** When switching IOAs that control disk units, ensure that all disk units that belong to that specific IOA are first removed from the auxiliary storage pool and are in a unconfigured status.

Virtual OptiConnect for IBM i logical partitions

The virtual OptiConnect feature provides high-speed interpartition communication within a managed system. The Virtual OptiConnect feature emulates external OptiConnect hardware by providing a virtual bus between logical partitions.

The virtual OptiConnect feature can be used only for communications between IBM i logical partitions. If you must enable communications with AIX or Linux logical partitions, use virtual Ethernet instead of the virtual OptiConnect feature.

To use the virtual OptiConnect feature on a logical partition, you must install OptiConnect for IBM i (a priced optional feature) on each IBM i logical partition that is to use virtual OptiConnect. If you use the Hardware Management Console (HMC) to create logical partitions on your managed system, you must also check the partition profile properties for each IBM i logical partition that is to use virtual OptiConnect and ensure that the **Use virtual OptiConnect** option is selected on the **OptiConnect** tab.

You can use the virtual OptiConnect feature without any additional hardware requirements.

Related information

OptiConnect

Expansion unit

You can add expansion units to many of the models to support additional features and devices. If you want to create logical partitions on your server, you must add an expansion unit that contains the additional hardware that you need for each logical partition.

Some expansion units can support only disk units (storage expansion unit), while others can support various hardware (system expansion unit). Expansion units generally contain one or more system I/O buses with various I/O devices.

If you assign more than 144 I/O slots to an AIX or Linux logical partition, ensure that the boot device for the logical partition is within the first 144 slots that are assigned to the logical partition. Also ensure that any PCIe3 2-Port 40GbE NIC RoCE QSFP+ Adapter (FC EC3A or FC EC3B) is within the first 144 slots. You can view the devices that are assigned to the first 144 slots of a logical partition by viewing the partition properties of the logical partition. Select the **Hardware** tab, then select the **I/O** tab, and then click the **Bus** column of the table to sort the devices in ascending order.

5250 CPW for IBM i logical partitions

5250 commercial processing workload (5250 CPW) is the capacity to perform 5250 online transaction processing (5250 OLTP) tasks on IBM i logical partitions.

A *5250 OLTP task* is a task that uses the 5250 data stream. Examples of 5250 OLTP tasks include the following:

- Any form of 5250 emulation, including Hardware Management Console (HMC) 5250, IBM Host On-Demand, IBM Personal Communications, and the 5250 emulation in the IBM i Access for Windows, Web, and Linux products
- 5250 Telnet or 5250 Display Station Pass-Through (DSPT) workstations
- Screen scrapers
- Interactive system monitors

You can use the IBM WebFacing tool to convert your 5250 OLTP applications into web-based applications that no longer need to use the 5250 data stream.

Application support for Linux logical partitions

Learn how to integrate Linux with IBM i applications and data.

Samba support with IBM i NetServer

Server Message Block (SMB) is a file-sharing protocol that is commonly used by Windows PCs. Whenever a network drive is mapped from a Windows PC to another Windows PC, the SMB TCP/IP protocol is being used.

Samba implements the SMB/CIFS standard on UNIX operating systems. This protocol enables file sharing among SMB-enabled operating systems, including IBM i with NetServer.

Samba allows Linux PCs and servers to interact with existing Windows PCs and file servers without requiring any additional software. IBM i NetServer supports Linux Samba clients.

You can use a Samba server to run printers and authenticate users, share files, and directories, just like Microsoft Windows PCs. Samba can also act as a Primary Domain Controller (PDC) or as a Backup Domain Controller (BDC) in your Windows network. You can use it to run OpenLDAP and add LDAP function to your Windows Network without the expense. You can use Samba and NetServer to share printers and files on IBM Power Systems or Linux partitions.

Accessing IBM i data by using Linux ODBC driver

The IBM i Open Database Connectivity (ODBC) driver for Linux allows you to access the IBM i database data from Linux applications that are written to the ODBC API. It is based on the ODBC driver in the IBM i Access for Windows product.

Related information

System i Access for Linux Open Database Connectivity

Examples: Logically partitioned systems

You can use the logical partitioning examples to consolidate servers, use computing resources more efficiently, and increase the flexibility of your enterprise.

Creating multiple client environments

You provide high-availability e-commerce services to a number of clients. You provide computing resources, applications, and technical support to each client, and each client can independently configure and use the applications that are running on the computing resources that you provide. In such an environment, it is essential to isolate the clients so that the clients can access only their resources. However, dedicating a physical server to each client is cost prohibitive, and does not allow you to easily increase or decrease the amount of computing resources that are used by each client.

Therefore, you decide to create a logical partition for each client. You install an operating system and applications on each logical partition. You can then use dynamic partitioning to add resources to logical partitions or remove resources from logical partitions as needed. If a client stops using your service, you can delete the logical partition for that client and reassign the resources to other logical partitions.

Testing new applications

You are a furniture manufacturer that uses an application to track inventory at your plant. A new version of the application is now available. You want to test this new version before you use it on your production server, but you do not have any money to buy separate test hardware.

Therefore, you decide to create a separate test environment on your managed system. You remove resources from the existing production environment, and you create a new logical partition that contains the resources that you removed from the production environment. You install an operating system and the new version of the inventory application on the logical partition. You can then use dynamic partitioning to move resources from the test logical partition to the production logical partition during peak production demand, and then return the resources to the test logical partition during testing. When you finish testing, you can delete the test logical partition, add the resources back to the production logical partition, and install the new version of the inventory application on the production system.

Integrating new acquisitions

You acquired a new company. Your new acquisition does not use the same applications for payroll, inventory, and billing that you do. You plan to consolidate your two companies onto a single set of applications, but it takes time to implement this consolidation. In the meantime, you are under pressure to reduce data center costs quickly.

Therefore, you decide to create logical partitions for the applications that are used by your new acquisition. You install an operating system and the applications that are used by the new company on the logical partition. If the combined workloads require more resources, you can use Capacity Upgrade on Demand (CUoD) to add processors and memory to the managed system, and then use dynamic partitioning to add these resources to the logical partitions. This solution allows you to save hardware costs immediately while you determine the best way to consolidate onto a single set of applications.

Planning for logical partitions

You can create logical partitions to distribute resources within a single server and make it function as if it were two or more independent servers. Before you create logical partitions, you must assess your current and future needs. You can then use this information to determine the hardware configuration that will meet your current needs and serve as a base for meeting your future needs.

Planning for logical partitions is a multi step process. The following tasks are recommended for planning for logical partitions.

_ **Assess your needs**

– Compile a list of the questions that you must answer before you create logical partition on an existing system or place your order for new hardware. The following is the list of questions:

- What are your existing workloads? How many resources do these workloads currently require (during typical usage and at peak usage)?
- What are your future needs? How will your existing workloads grow over the life of your system? How many new workloads do you have to support over the life of your system?
- Do you have an existing system onto which you can consolidate the workloads? Must you upgrade the existing system before you consolidate the workloads? Does it make more sense to purchase a new system for these workloads?
- What physical infrastructure will you have to support any new hardware? Can your current location accommodate the new hardware? Must you upgrade your power infrastructure or your cooling infrastructure?
- Will your new hardware work with your existing hardware?
- Which hardware features will you use? For example, do you want to use virtual I/O to consolidate I/O resources? Must you obtain activation codes or enablement codes to use these features?
- Must you obtain more licenses to run your applications? If so, how many more licenses do you need?
- Does the support strategy for your new hardware differ from the support strategy for your existing hardware? If so, what changes must you make to maximize the effectiveness of the new support strategy?
- Must you migrate your workloads onto new hardware? If so, what must you do to migrate these workloads?

_ **Learn about your system and its features**

– Your system has many features that allow you to use system resources more efficiently and simplify daily tasks. For more information about what these features are and how these features work, see "Logical partition overview" on page 2.

_ **Learn about planning tools**

– IBM provides the following tools that you can use to assess your needs, determine the hardware that you need to accommodate existing and future needs, and compile an order for the hardware that you need:

IBM Prerequisite website
The IBM Prerequisite website provides you with compatibility information for hardware features. This site helps you plan a successful system upgrade by providing you with the prerequisite information for features that you currently have or plan to add to your system.

IBM Systems Workload Estimator
The IBM Systems Workload Estimator (WLE) estimates the computer resources that are required for Domino®, WebSphere® Commerce, WebSphere, Web Serving, and traditional workloads. The WLE projects the most current server models that meet the capacity requirements that are within the CPU percent utilization objectives.

AIX Performance Toolbox for POWER®
The AIX Performance Toolbox (PTX) for POWER is a licensed program that provides a comprehensive tool for monitoring and tuning system performance in distributed environments.

_ **Take inventory of your current environment**

– Monitor resource usage on your existing servers to determine the amounts of resources that you currently use in your operation. You will use this information as a basis for determining the resources that you require on the consolidated system. The Performance Monitor (PM) information that you gather from your existing systems gives you the information that you need to analyze existing workloads.

_ **Perform capacity planning**

– Analyze the workloads that are to be consolidated onto your managed system and determine the amounts of resources that these workloads require. You will also want to calculate the resources that you will need for future growth and determine whether your hardware can accommodate this growth. To analyze your current workloads, use your PM information as input for the WLE. The WLE uses this input to determine the resources that you need for the consolidated workloads. The WLE also allows you to project how many resources you will need in the future.

_ **Decide which tool you want to use to create logical partitions and manage the system**

– Determine whether you want to use the Hardware Management Console (HMC), or the Virtual Partition Manager to create logical partitions and manage the system. To learn about these tools, see "Logical partitioning tools" on page 10.

_ **Decide if you want your operating systems to share I/O resources with each other**

– Determine whether you want to set your logical partitions to use virtual I/O resources from a Virtual I/O Server logical partition. For more information, see Virtual I/O Server.

_ **Design and validate your logical partition configuration**

– Design the logical partitions that you will create on the managed system, and assign resources to each logical partition so that the logical partitions can perform their assigned tasks efficiently.

_ **Design network infrastructure to connect logical partitions with each other and with external**
_ **networks**

Determine what types of physical and virtual adapters you want to use to connect logical partitions to each other and to external networks. For more information about the different methods that you can use to connect logical partitions with each other and with external networks, see “I/O devices” on page 46.

_ **Identify how the managed system communicates with the HMC**

– Determine how you want to connect your managed system and its logical partitions with the HMC that manages the system. For more information about the ways in which you can connect your managed system with the HMC, see HMC network connections.

_ **Determine a service and support strategy**

– Determine how to apply fixes to your server and identify problems that need to be reported to your service provider. The HMC can be configured to report most problems to your service provider automatically. For more information about how to set up the HMC to report problems, see Configuring the local console to report errors to service and support.

_ **Plan for software licensing in a partitioned environment**

– Determine the number of software licenses that you need for your logical partition configuration. For instructions, see “Software licensing for IBM licensed programs on logical partitions” on page 75.

Trusted Boot

Trusted Boot is a feature of Power® Security and Compliance (PowerSC). Trusted Boot uses the Virtual Trusted Platform Module (VTPM) as described by the Trusted Computing Group. Up to 60 logical partitions per server can be configured to have their own unique VTPM by using the Hardware Management Console (HMC). The VTPM is used to record the system boot and in association with the AIX Trusted Execution technology, provides security and assurance of the boot image on disk, on the entire operating system, and in the application layers.

The VTPM is a software implementation of the Trusted Platform Module (TPM) specification, as described by the Trusted Computing Group. The Trusted Platform Module is implemented as a physical chip on computer systems.

You can create a VTPM as part of the initial logical partitioning (by using the HMC Partitioning wizard), or you can dynamically enable the device. When dynamically enabled, the VTPM becomes active only when the logical partition is restarted.

The VTPM enables the AIX environment of the logical partition to use Trusted Boot capability. When a VTPM is associated with a logical partition, being booted, components of the boot take cryptographic hashes of relevant data and of components that can be run in the future, for example the AIX boot loader. These cryptographic hashes are securely copied to storage that is controlled by the VTPM. After the logical partition is operational, other users can then securely retrieve the hashes by using a process that is known as remote attestation. The hashes can then be examined to determine whether the logical partition booted in a trusted configuration so that users can take action if required.

To use a VTPM, the logical partition must have the following resources:

- The maximum memory setting of the logical partition must be greater than of 1 GB for the active profile.
- Each VTPM requires permanent storage for the lifetime of the device. A normal logical partition uses 6 KB of system nonvolatile RAM. This storage requirement imposes a limitation on the number of VTPMs per server.

Permanent data that is stored by the VTPM contains sensitive information about the trust of the VTPM feature. For example, the first time each VTPM is operated a public-private key pair that is known as the Endorsement Key (EK) is generated and then permanently stored. This action allows the VTPM to be identified by other users during the lifetime of the device. The permanent data, including the EK, is deleted when the VTPM device is removed by the console.

To maintain the sensitivity of the stored data, the data is secured by the trusted system key, which is under the control of the HMC. The trusted system key secures the VTPM data but has an impact on logical partition mobility, and suspend features for logical partitions that are enabled for VTPM. A logical partition that is enabled for VTPM must adhere to the following prerequisites to support the logical partition mobility, and suspend features:

- To migrate a logical partition with VTPM enabled, both systems must have the same trusted system key.
- To successfully change the trusted system key, no logical partition with VTPM enabled can be in the suspend state. The HMC cannot change the key until suspended logical partitions with VTPM enabled, are resumed or powered off.

Trusted Firewall

With Virtual I/O Server (VIOS) Version 2.2.1.4 or later, and POWER9 processor-based servers with firmware at level FW740, or later, you can use the Trusted Firewall feature. Trusted Firewall is a feature of the PowerSC Editions. You can use the Trusted Firewall feature to provide a virtual firewall that allows network filtering and control within the local server. The virtual firewall improves performance and reduces the consumption of network resources by allowing direct and secure network traffic between logical partitions that are on different VLANs of the same server.

With the Trusted Firewall feature, you can perform LAN routing functions between logical partitions on the same server by using the Security Virtual Machine (SVM) kernel extension. By using the Trusted Firewall feature, logical partitions that are on different virtual LANs of the same server can communicate by using the shared Ethernet adapter (SEA). Trusted Firewall is supported on AIX, IBM i, and Linux logical partitions.

Preparing to configure Active Memory Expansion

Before you configure Active Memory Expansion for a logical partition, you need to ensure that your system meets the configuration requirements. Optionally, you can run the Active Memory Expansion planning tool.

About this task

To prepare to configure Active Memory Expansion for a logical partition, complete the following tasks:

Procedure

1. Ensure that your system meets the following configuration requirements:
   - The server on which the logical partition runs is a POWER7, or later.
   - The AIX operating system that runs in the logical partition is at version 6.1 with Technology Level 4 and Service Pack 2, or later.
   - The HMC that manages the server is at version 7, release 7.1.0, or later.
2. Optional: Run the Active Memory Expansion planning tool, which is the **amepat** command, from the AIX command-line interface.

   The planning tool monitors your current workload for a specified amount of time and generates a report. The report provides the following information:

   - Several configuration possibilities for Active Memory Expansion on the logical partition.
   - Recommendation for an initial configuration for Active Memory Expansion on the logical partition.

   For each configuration possibility, and the recommended configuration, the planning tool provides the following configuration information:

   - The amount of memory to assign to the logical partition.
   - The amount of additional processing resources to assign to the logical partition.

- The expansion factor to set for the logical partition.
- The amount of memory that you save by configuring Active Memory Expansion on the logical partition. This statistic can help you determine whether Active Memory Expansion is right for your workload. Some workloads benefit more from Active Memory Expansion than others.

What to do next

After you prepare to configure Active Memory Expansion, you can enable Active Memory Expansion on the server by entering the activation code.

Configuration requirements for shared memory

Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

System requirements

- The server must be a POWER7 processor-based server, or later.
- The server firmware must be at release 3.4.2, or later.
- The Hardware Management Console (HMC) must be at version 7 release 3.4.2, or later.
- The PowerVM Active Memory Sharing technology must be activated. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

Paging VIOS partition requirements

- VIOS partitions that provide access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool (hereafter referred to as *paging VIOS partitions*) cannot use shared memory. Paging VIOS partitions must use dedicated memory.
- Paging VIOS partitions must be at version 2.1.1, or later.
- On HMC-managed systems, consider configuring separate VIOS partitions as server partitions and paging VIOS partitions. For example, configure one VIOS partition to provide virtual resources to the shared memory partitions. Then, configure another VIOS partition as a paging VIOS partition.
- On HMC-managed systems, you can configure multiple VIOS partitions to provide access to paging space devices. However, you can assign only up to two of those VIOS partitions to the shared memory pool at any time.

Requirements for shared memory partitions

- Shared memory partitions must use shared processors.
- You can assign only virtual adapters to shared memory partitions. This means that you can dynamically add only virtual adapters to shared memory partitions. More specifically, the following table lists the virtual adapters that you can assign shared memory partitions.

Table 12. Virtual adapters that you can assign to shared memory partitions

AIX and Linux shared memory partitions	**IBM i shared memory partitions**
– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial adapters	– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial server adapters

Table 13. Virtual adapters that you can assign to shared memory partitions

Linux shared memory partitions
– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial adapters

You cannot assign Host Ethernet Adapters (HEA) or host connection adapters (HCA) to shared memory partitions.

- Shared memory partitions cannot use the barrier synchronization register.
- Shared memory partitions cannot use huge pages.
- AIX must be at version 6.1 Technology Level 3, or later, to run in a shared memory partition.
- IBM i must be at 6.1 with PTF SI32798, or later, to run in a shared memory partition.
- Virtual OptiConnect must not be enabled on IBM i shared memory partitions.
- SUSE Linux Enterprise Server must be at version 11, or later, to run in a shared memory partition.
- Red Hat Enterprise Server Version 6, or later, to run in a shared memory partition.
- You cannot configure IBM i logical partitions that provide virtual resources to other logical partitions as shared memory partitions. Logical partitions that provide virtual resources to other logical partitions in a shared memory environment must be VIOS partitions.

Requirements for paging space devices

- The paging space devices for AIX or Linux shared memory partitions must be at least the size of the maximum logical memory of the shared memory partition.
- The paging space devices for IBM i shared memory partitions must be at least the size of the maximum logical memory of the shared memory partition plus 8 KB for every megabyte. For example, if the maximum logical memory of the shared memory partition is 16 GB, its paging space device must be at least 16.125 GB.
- Paging space devices can be assigned only to one shared memory pool at a time. You cannot assign the same paging space device to a shared memory pool on one system and to another shared memory pool on another system at the same time.
- Paging space devices that are accessed by a single paging VIOS partition must meet the following requirements:
  - They can be physical or logical volumes.
  - They can be located in physical storage on the server or on a storage area network (SAN).
- Paging space devices that are accessed redundantly by two paging VIOS partitions must meet the following requirements:
  - They must be physical volumes.
  - They must be located on a SAN.
  - They must be configured with global IDs.
  - They must be accessible to both paging VIOS partitions.
  - The reserve attribute must be set to no reserve. (The VIOS automatically sets the reserve attribute to no reserve when you add the paging space device to the shared memory pool.)
- Physical volumes that are configured as paging space devices cannot belong to a volume group, such as the `rootvg` volume group.
- Logical volumes that are configured as paging space devices must be located in a volume group that is dedicated for paging space devices.

- Paging space devices must be available. You cannot use the physical volume or logical volume as a paging space device if it is already configured as a paging space device or virtual disk for another logical partition.
- Paging space devices cannot be used to boot a logical partition.
- After you assign a paging space device to the shared memory pool, you must manage the device by using the **Create/Modify Shared Memory Pool** wizard on the HMC. Do not change or remove the device by using other management tools.

Related concepts

Paging space devices on systems that are managed by an HMC
Learn about the location requirements, size requirements, and redundancy preferences for paging space devices on systems that are managed by a Hardware Management Console (HMC).

Configuration requirements and restrictions for the remote restart capability of a logical partition

An AIX, Linux, or IBM i logical partition that supports the remote restart feature must have its configuration information and persistent data stored external to the server on persistent storage.

With the Hardware Management Console (HMC) Version 7.6.0, or later, a logical partition can be enabled for remote restart on any server that supports the remote restart capability. To perform a remote restart operation by using the HMC command line, the HMC must be at version 8.1.0 or later. A logical partition can recover from a server outage by starting on another server.

The following are the configuration requirements for a logical partition with the remote restart capability:

- When the HMC is at Version 8.2.0, or later, and when you select the PowerVM simplified remote restart feature, you need not assign a reserved storage device from the reserved storage device pool to the logical partition before the logical partition is activated.
- When the Hardware Management Console (HMC) is at Version 8.4.0, or later, and the Virtual I/O Server (VIOS) at Version 2.2.4.0, or later, the remote restart of logical partitions that use shared storage pool devices is supported.
- When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can remotely restart a logical partition that has migratable single root I/O virtualization (SR-IOV) logical ports.

  Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.
- When the HMC is at Version 9.2.950, or later, and the firmware is at a level FW950, or later, for the remote restart operation to be successful on the logical partition that supports platform keystore capability, the user-defined system key that is configured on both the source and the destination system must match.

Restriction:

The following are the restrictions for a logical partition that supports the remote restart feature:

- The logical partition must not have physical I/O adapters.
- If you are using HMC Version 9.1.930, or earlier, the logical partition must not have SR-IOV logical ports.
- The logical partition must not be a full system partition, or a Virtual I/O Server (VIOS) partition.
- The logical partition must not be enabled for redundant error path reporting.
- The logical partition must not have a barrier-synchronization register (BSR).
- The logical partition must not have huge pages.

- The logical partition must not have its `rootvg` volume group on a logical volume or have any exported optical devices.
- The logical partition must not be set to automatically start with the server.
- The logical partition must not have any Host Channel Adapter resources.
- The logical partition must not be set as the service partition for the server.
- The logical partition must not have any Host Ethernet Adapter resources.
- The logical partition must not belong to a workload group.
- The logical partition must not use shared memory.
- The logical partition must not have a Virtual Trusted Platform Module (VTPM) enabled.
- The logical partition must not use Virtual Station Interface.

In addition to the restrictions for a logical partition that supports the remote restart feature, IBM i partitions have some additional restrictions.

Restriction:

The following list illustrates the restrictions for an IBM i logical partition that supports the remote restart feature:

- The logical partition must not have a virtual server SCSI adapter. Also, the logical partition must not have a virtual SCSI client adapter that is associated with a server adapter that is not on a VIOS partition.
- The logical partition must not have HSL (High-Speed Link) OptiConnect or Virtual OptiConnect enabled.

Verifying that the server supports partitions that are capable of the simplified version of the remote restart feature

Before planning to enable the simplified version of the remote restart capability of a logical partition, verify that the server supports partitions that are capable of the simplified version of the remote restart feature by using the Hardware Management Console (HMC).

About this task

To verify that the server supports partitions that are capable of the simplified version of the remote restart feature, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Click **Licensed Capabilities**. The Licensed Capabilities page lists the features that are supported by the server.
5. In the Licensed Capabilities page, verify the list of features displayed.
   - If **PowerVM Partition Simplified Remote Restart Capable** is marked by the ✓ icon, server supports the simplified version of the remote restart feature.
   - If **PowerVM Partition Simplified Remote Restart Capable** is marked by the -- icon, the server does not support the simplified version of the remote restart feature.
6. Click **OK**.

Verifying that the logical partition supports the simplified version of the remote restart feature

You can use the Hardware Management Console (HMC) to verify whether the logical partition supports the simplified version of the remote restart feature.

About this task

To verify that the logical partition supports the simplified version of the remote restart feature, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **View Partition Properties**.
4. Click the **General Properties** tab.
    - If the **Simplified Remote Restart** check box is selected, the logical partition supports the simplified version of the remote restart feature.
    - If the **Simplified Remote Restart** check box is not selected, the logical partition does not support the simplified version of the remote restart feature.
5. Click **OK**.

Verifying that the server supports Virtual Trusted Platform Module

Before planning to enable a Virtual Trusted Platform Module (VTPM) on a logical partition, verify that the server supports VTPM by using the Hardware Management Console (HMC).

About this task

To verify that the server supports partitions that are capable of VTPM, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
3. Click **Advanced**. The server supports VTPM if you can view information about VTPM.
4. Click **OK**.

Enabling the platform keystore capability on a logical partition

With the HMC version 9.2.950, you can enable the platform key store capability if the system supports the feature.

Before enabling the platform keystore feature on a logical partition, verify that the IBM Power system server supports the platform keystore capability by using the **lssyscfg** command along with **-F** attribute. Before running this command, ensure that the system is in a stand by or operating state. If the system supports the platform keystore capability, you can also use the **lssyscfg** command to view the range of the keystore size.

When the HMC is at V9.2.950.0, or later, and when the firmware is at level FW950, or later, you can choose the keystore size for a logical partition. You can choose the value for the keystore size within the range supported by the system. The value of 0 kilobytes (KB) indicates that the keystore function is disabled for the logical partition. To view the minimum and maximum value for the keystore size, type the following the command:

```
lssyscfg -r sys -m <system_name> -Flpar_keystore_min_kbytes, lpar_keystore_max_kbytes
```

Related information

lssyscfg command

chtskey command

Configuring the Virtual I/O Server for the VSN capability

If you are using the Hardware Management Console (HMC) Version 7 Release 7.7.0, or later, you can use Virtual Station Interface (VSI) profiles with virtual Ethernet adapters in logical partitions and assign the Virtual Ethernet Port Aggregator (VEPA) switching mode to virtual Ethernet switches.

When you use the Virtual Ethernet Bridge (VEB) switching mode in virtual Ethernet switches, the traffic between logical partitions is not visible to the external switches. However, when you use the VEPA switching mode, the traffic between logical partitions is visible to the external switches. This visibility helps you to use features such as security that are supported by the advanced switching technology. Automated VSI discovery and configuration with the external Ethernet bridges simplifies the switch configuration for the virtual interfaces that are created with logical partitions. The profile-based VSI management policy definition provides flexibility during configuration and maximizes the benefits of automation.

The configuration requirements on the Virtual I/O Server (VIOS) to use the VSN capability follow:

- At least one VIOS logical partition that is servicing the virtual switch must be active and must support the VEPA switching mode.
- The external switches that are connected to the shared Ethernet adapter must support the VEPA switching mode.
- The **`lldp`** daemon must be running on the VIOS and must be managing the shared Ethernet adapter.
- From the VIOS command-line interface, run the **chdev** command to change the value of the *lldpsvc* attribute of the shared Ethernet adapter device to *yes*. The default value of the *lldpsvc* attribute is *no*. Run the **`lldpsync`** command to notify the change to the running **`lldpd`** daemon.

  Note: The *lldpsvc* attribute must be set to the default value before you remove the shared Ethernet adapter. Otherwise, removal of the shared Ethernet adapter fails.
- For redundancy shared Ethernet adapter setup, the trunk adapters might be attached to a virtual switch that is set to the VEPA mode. In this case, attach the control channel adapters of the shared Ethernet adapter to another virtual switch that is always set to the virtual Ethernet bridging (VEB) mode. The shared Ethernet adapter that is in the high availability mode does not work when the control channel adapter that is associated with the virtual switches is in the VEPA mode.

Restriction: To use VSN capability, you cannot configure a shared Ethernet adapter to use link aggregation or an Etherchannel device as the physical adapter.

Verifying that the server uses the virtual server network

Before planning to enable the virtual server network (VSN), verify that the server uses VSN by using the Hardware Management Console (HMC).

About this task

As of HMC Version 7 Release 7.7.0, you can assign the Virtual Ethernet Port Aggregator (VEPA) switching mode to virtual Ethernet switches that are used by the virtual Ethernet adapters of the logical partitions. The VEPA switching mode uses the features that are supported by the advanced virtual Ethernet switch technology. A logical partition whose virtual Ethernet adapters use virtual switches that are enabled with the VEPA switching mode, uses VSN.

You can use the **`lssyscfg`** command to verify that the server uses VSN.

Verifying that the server supports single root I/O virtualization

Before you enable single root I/O virtualization (SR-IOV) shared mode for an SR-IOV capable adapter, verify that the server supports the SR-IOV feature by using the Hardware Management Console (HMC). SR-IOV is a Peripheral Component Interconnect Special Interest Group specification to allow multiple partitions that are running simultaneously within a single computer to share a Peripheral Component Interconnect-Express (PCIe) device.

About this task

To verify that the server supports SR-IOV, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Click **Licensed Capabilities**. The Licensed Capabilities page lists the features that are supported by the server.
5. In the Licensed Capabilities page, verify the list of features displayed.
   - If **SR-IOV Capable** is marked by the ✓ icon, the SR-IOV adapter can be configured in the shared mode and can be shared by multiple logical partitions.
   - If **SR-IOV Capable** is marked by the -- icon, the SR-IOV adapter can be configured in the shared mode but can be used by only one logical partition.
   - If **SR-IOV Capable** is not displayed, the server does not support the SR-IOV feature.
6. Click **OK**.

Verifying the logical port limit and the owner of the SR-IOV adapter

You can view the logical port limit and the owner of the single root I/O virtualization (SR-IOV) adapter by using the Hardware Management Console (HMC).

About this task

To view the logical port limit and the owner of the SR-IOV adapter, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Click **Licensed Capabilities**. The Licensed Capabilities page lists the features that are supported by the server.
5. In the **Properties** area, click the **Processor, Memory, I/O** tab. In the **Physical I/O Adapters** area, the table displays the **SR-IOV capable (Logical Port Limit)** and the **Owner** details about the SR-IOV adapter.

- The **SR-IOV capable (Logical Port Limit)** column displays whether the slot or the adapter is SR-IOV capable, and the maximum number of logical ports this slot or the adapter can support. If the slot or the adapter is SR-IOV capable but is assigned to a partition, the **SR-IOV capable (Logical Port Limit)** column indicates that the slot or the adapter is in the dedicated mode.

- The **Owner** column displays the name of the current owner the physical I/O. The value of this column can be any of the following values:
  - When an SR-IOV adapter is in the shared mode, **Hypervisor** is displayed in this column.
  - When an SR-IOV adapter is in the dedicated mode, **Unassigned** is displayed when the adapter is not assigned to any partition as a dedicated physical I/O.
  - When an SR-IOV adapter is in the dedicated mode, the logical partition name is displayed when the adapter is assigned to any logical partition as a dedicated physical I/O.

Verifying that the server supports IBM i native I/O capability

You can verify that the server supports IBM i native I/O capability by running the **`lssyscfg`** command that is available in the Hardware Management Console (HMC) command-line interface.

Procedure

To verify whether the server supports IBM i native I/O capability, type the following command from the HMC command line:

```
lssyscfg -r sys -m <server> -F os400_native_io_capable
```

where *server* is the user-defined name of the server you want to check. The *os400_native_io_capable* attribute can have any one of the following values:

- unavailable - When the server firmware is at a level that is earlier than FW860. When this value is returned, the HMC does not know whether the server supports IBM i native I/O capability.
- 0 - When the server does not support IBM i native I/O capability.
- 1 - When the server supports IBM i native I/O capability.

Preparing to configure shared memory

Before you configure the shared memory pool and create logical partitions that use shared memory (hereafter referred to as *shared memory partitions*), you need to plan for the shared memory pool, the shared memory partitions, the paging space devices, and the Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*).

Related concepts

Paging space devices on systems that are managed by an HMC
Learn about the location requirements, size requirements, and redundancy preferences for paging space devices on systems that are managed by a Hardware Management Console (HMC).

Logical memory
Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Preparing to configure shared memory on a system that is managed by an HMC

Before you configure the shared memory pool and create logical partitions that use shared memory, you need to determine the size of the shared memory pool, the amount of memory to assign to each shared memory partition, the number of paging space devices to assign to the shared memory pool, and the redundancy configuration of the Virtual I/O Server logical partitions that you assign to the shared memory pool.

Before you begin

Before you start, verify that your system meets the requirements for configuring shared memory. For instructions, see “Configuration requirements for shared memory” on page 65.

About this task

To prepare to configure the shared memory pool and shared memory partitions, complete the following steps:

Procedure

1. Assess your needs, take inventory of your current environment, and plan for capacity. For instructions, see "Planning for logical partitions" on page 61. Determine the following information:
   a) Determine the number of shared memory partitions to assign to the shared memory pool.
   b) Determine the amount of logical memory to assign as the allocated, minimum, and maximum logical memory for each shared memory partition.

      You can apply the same general guidelines that you might use to assign the allocated, minimum, and maximum dedicated memory to logical partitions that use dedicated memory. For example:

      - Do not assign the maximum logical memory to a value that is higher than the amount of logical memory that you plan to dynamically add to the shared memory partition.
      - Set the minimum logical memory to a value that is high enough for the shared memory partition to successfully activate.
2. Determine the amount of physical memory to assign to the shared memory pool.

   For instructions, see "Determining the size of the shared memory pool" on page 74.
3. Prepare for paging space devices:
   a) Determine the number of paging space devices to assign to the shared memory pool.

      The HMC assigns one paging space device to each shared memory partition that is active. Thus, the fewest number of paging space devices that must be assigned to the shared memory pool is equal to the number of shared memory partitions that you plan to run simultaneously. For example, you assign 10 shared memory partitions to the shared memory pool and you plan to run eight of the shared memory partitions simultaneously. Thus, you assign at least eight paging space devices to the shared memory pool.
   b) Determine the size of each paging space device:

      - For AIX and Linux Linux shared memory partitions, the paging space device must be at least the size of the maximum logical memory size of the shared memory partition that you identified in step 1b. For example, you plan to create an AIX shared memory partition with a maximum logical memory size of 16 GB. The paging space device must be at least 16 GB.
      - For IBM i shared memory partitions, the paging space device must be the size of the maximum logical memory size of the shared memory partition that you identified in step 1b multiplied by 129/128. For example, you plan to create an IBM i shared memory partition with a maximum logical memory size of 16 GB. The paging space device must be at least 16.125 GB.
      - Consider creating paging space devices that are large enough to be used by shared memory partitions with multiple partition profiles.
   c) Determine whether each paging space device resides in physical storage on the server or on a storage area network (SAN).

      Paging space devices that are accessed by a single Virtual I/O Server (VIOS) logical partition can be located in physical storage on the server or on a SAN. Paging space devices that are redundantly accessed by two paging VIOS partitions must be located on a SAN.
4. Prepare for paging VIOS partitions:
   a) Determine which Virtual I/O Server (VIOS) logical partitions can be assigned to the shared memory pool as paging VIOS partitions.

      A paging VIOS partition provides access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool. A paging VIOS partition can be any active Virtual I/O Server (version 2.1.1, or later) that has access to the paging space devices that you plan to assign to the shared memory pool.
   b) Determine the number of paging VIOS partitions to assign to the shared memory pool.

You can assign 1 or 2 paging VIOS partitions to the shared memory pool:

- When you assign a single paging VIOS partition to the shared memory pool, it must have access to all of the paging space devices that you plan to assign to the shared memory pool.
- When you assign two paging VIOS partitions to the shared memory pool, each paging space device that you plan to assign to the shared memory pool must be accessible to at least one paging VIOS partition. However, usually when you assign two paging VIOS partitions to the shared memory pool, they redundantly access one or more paging space devices.

c) If you plan to assign two paging VIOS partitions to the shared memory pool, determine how you want to configure redundancy for the shared memory partitions:

i) Determine which shared memory partitions to configure to use redundant paging VIOS partitions. For each shared memory partition, this means that both paging VIOS partitions can access the shared memory partition's paging space device.

ii) Determine which paging VIOS partition to assign as the primary paging VIOS partition and which paging VIOS partition to assign as the secondary paging VIOS partition for each shared memory partition. The hypervisor uses the primary paging VIOS partition to access the paging space device that is assigned to the shared memory partition. If the primary VIOS partition becomes unavailable, the hypervisor uses the secondary paging VIOS partition to access the paging space device that is assigned to the shared memory partition.

5. Determine the number of additional processor resources that are needed for the paging VIOS partitions.

   To read and write data between the paging space devices and the shared memory pool, the paging VIOS partitions require more processing resources. The amount of additional processing resources that are needed depends on the frequency that the paging VIOS partition reads and writes the data. The more frequently that the paging VIOS partition reads and writes the data, the more frequently the paging VIOS partition performs I/O operations. More I/O operations require more processing power. In general, the frequency that the paging VIOS partition reads and writes data can be affected by the following factors:

   - The degree to which the shared memory partitions are overcommitted. In general, greatly overcommitted shared memory partitions require the paging VIOS partition to read and write data more often than slightly overcommitted shared memory partitions.
   - The I/O rates of the storage subsystem on which the paging space devices are located. In general, paging space devices with faster I/O rates (such as a SAN) enable the paging VIOS partition to read and write data more often than paging space devices with slower I/O rates (such as storage in the server).

   You can use the IBM Systems Workload Estimator (WLE) to determine the number of processor resources that are needed for paging VIOS partitions.

Related concepts

Performance considerations for shared memory partitions
You can learn about performance factors (such as shared memory overcommitment) that influence the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). You can also use shared memory statistics to help you determine how to adjust the configuration of a shared memory partition to improve its performance.

Determining the size of the shared memory pool

You need to consider the degree to which you want to over-commit the physical memory in the shared memory pool, the performance of the workloads when they are running in a shared memory configuration that is over-committed, and the minimum and maximum boundaries of the shared memory pool.

About this task

To determine the size of the shared memory pool, consider the following factors:

Procedure

1. Consider the degree to which you want to over-commit the physical memory in the shared memory pool.

    - When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.
    - When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

2. Consider the performance of the workloads when running in a shared memory configuration that is over-committed. Some workloads perform well in a shared memory configuration that is logically over-committed, and some workloads can perform well in a shared memory configuration that is physically over-committed.

    Tip: In general, more workloads perform better in logically over-committed configurations than physically over-committed configurations. Consider limiting the degree to which you physically over-commit the shared memory pool.

3. The shared memory pool must be large enough to meet the following requirements:

    a) The shared memory pool must be large enough to provide each shared memory partition with its I/O entitled memory when all of the shared memory partitions are active. When you create a shared memory partition, the Hardware Management Console (HMC) automatically determines the I/O entitled memory for the shared memory partition.

    After you activate the shared memory partitions, you can view statistics about how the operating systems use their I/O entitled memory and adjust the I/O entitled memory of the shared memory partitions accordingly.

    b) A small portion of the physical memory in the shared memory pool is reserved for the hypervisor so that it can manage shared memory resources. The hypervisor requirement is a small amount of physical memory per shared memory partition + 256 MB.

    Tip: To ensure that you can successfully activate the shared memory partitions, assign at least the following amount of physical memory to the shared memory pool: (the sum of the minimum logical memory that is assigned to all of the shared memory partitions that you plan to run concurrently) + (the required 256 MB of reserved firmware memory).

4. When the shared memory pool is equal to or greater than the sum of the assigned logical memory of all the shared memory partitions plus the required amount of reserved firmware memory, the initial shared memory configuration is not over-committed. Therefore, the amount of physical memory that you assign to the shared memory pool need not exceed the sum of the assigned logical memory of all the shared memory partitions plus the required amount of reserved firmware memory.

Software licensing for IBM licensed programs on logical partitions

If you use IBM licensed programs such as AIX and IBM i on a server with logical partitions, consider how many software licenses are required for your logical partition configuration. Careful consideration of your software might help minimize the number of software licenses that you must purchase.

Software license behavior varies by software product. Each solution provider has its own licensing strategy. If you use licensed programs from solution providers other than IBM, consult the documentation from those solution providers to determine the licensing requirements for those licensed programs.

With some servers, you can purchase IBM i licenses on a per-user basis. For more information about IBM i licenses, see Working with software agreements and licenses in the IBM Knowledge Center.

Many IBM licensed programs allow you to purchase licenses based on the number of processors that the licensed program uses on a managed system as a whole. An advantage of this processor-based licensing method is that it allows you to create multiple logical partitions without having to purchase separate licenses for each logical partition. Also, this method caps the number of licenses that you need for a managed system. You need never obtain more licenses for a single licensed program than the number of processors on the managed system.

The main complicating factor in calculating the number of licenses that are required on a managed system with logical partitions that use processor-based licensing is the fact that a logical partition that uses uncapped shared processors can use up to its assigned number of virtual processors. When you use processor-based licensing, ensure that the number of virtual processors on uncapped logical partitions are set so that each IBM licensed program does not use more processors than the number of processor-based licenses that you purchased for that IBM licensed program.

The number of licenses that are required for a single IBM licensed program on a managed system that uses processor-based licensing is the **lesser** of the following two values:

- The total number of activated processors on the managed system.
- The maximum number of processors that can be used by the IBM licensed program on the managed system. The maximum number of processors that can be used by the IBM licensed program on the managed system is the **sum** of the following two values:
  - The total number of processors that are assigned to all logical partitions that use dedicated processors and that run the IBM licensed program.
  - The sum of the maximum number of processing units that can run the IBM licensed program on **each** shared processor pool, rounded up to the next integer. The maximum number of processing units that can run the IBM licensed program on each shared processor pool is the **lesser** of the following two values:
    - The total number of processing units that are assigned to capped logical partitions that run the IBM licensed program, plus the total number of virtual processors that are assigned to uncapped logical partitions that run the IBM licensed program.
    - The maximum number of processing units that are specified for the shared processor pool. (For the default shared processor pool, this number is the total number of activated processors on the managed system minus the total number of processors that are assigned to all logical partitions that use dedicated processors and that are not set to share processors with shared processor logical partitions. Use of Capacity on Demand (CoD) can increase the number of activated processors on the managed system, which can cause the managed system to go out of compliance if you do not allow for CoD use. Also, if there are logical partitions that use dedicated processors, that run the IBM licensed program, and that are set to share processors with shared processor logical partitions, then you can deduct the processors for these dedicated processor logical partitions from the maximum number of processing units for the default shared processor pool total because you already counted these dedicated processors in the dedicated processor logical partition total.)

When you use processor-based licensing, ensure that the managed system is in compliance with the license agreement for each IBM licensed program that is installed on the managed system. If you have a managed system that can use multiple shared processor pools, you can use the multiple shared processor pool feature of the Hardware Management Console (HMC) to ensure that your managed system remains in compliance with these license agreements. You can configure a shared processor pool with a maximum processing unit value equal to the number of licenses that you have for your managed system, and then set all logical partitions that use the IBM licensed program so that they use that shared processor pool. The logical partitions in the shared processor pool cannot use more processors than the maximum processing unit value that is set for the shared processor pool, so the managed system remains in compliance with the per-processor license agreement.

For example, Company Y has obtained three processor-based IBM i licenses for a managed system with four processors and four logical partitions. The managed system has only one shared processing pool, and all four logical partitions use the shared processor pool, so all four of the processors on the managed system are in the shared processor pool. The configuration of the logical partitions is as follows.

Table 14. Logical partition configuration in compliance with license agreement

Logical partition name	Operating system	Processing mode	Sharing mode	Processing units	Virtual processors	Maximum number of processors that can be used by the logical partition
Partition A	IBM i	Shared	Uncapped	1.75	2	2.00 (the number of virtual processors for the uncapped shared logical partition)
Partition B	IBM i	Shared	Capped	0.60	1	0.60 (the number of processing units for the capped shared logical partition)
Partition C	IBM i	Shared	Capped	0.40	1	0.40 (the number of processing units for the capped shared logical partition)
Partition D	Linux	Shared	Uncapped	1.25	2	2.00 (the number of virtual processors for the uncapped shared logical partition)

This configuration has three IBM i logical partitions and one Linux logical partition on the managed system. The three IBM i logical partitions can use a maximum of 3.00 processors (2.00 for Partition A, 0.60 for Partition B, and 0.40 for Partition C). The managed system has three IBM i licenses, so the managed system is in compliance with the IBM i license agreement.

For an example of a logical partition configuration that is out of compliance with a licensing agreement, the system administrator at Company Y changes the sharing mode of Partition B and Partition C from capped to uncapped. The following table shows the new logical partition configuration.

Table 15. Logical partition configuration out of compliance with license agreement (first example)

Logical partition name	Operating system	Processing mode	Sharing mode	Processing units	Virtual processors	Maximum number of processors that can be used by the logical partition
Partition A	IBM i	Shared	Uncapped	1.75	2	2.00 (the number of virtual processors for the uncapped shared logical partition)
Partition B	IBM i	Shared	Uncapped	0.60	1	1.00 (the number of virtual processors for the uncapped shared logical partition)
Partition C	IBM i	Shared	Uncapped	0.40	1	1.00 (the number of virtual processors for the uncapped shared logical partition)
Partition D	Linux	Shared	Uncapped	1.25	2	2.00 (the number of virtual processors for the uncapped shared logical partition)

In this configuration, the three IBM i logical partitions can use a maximum of 4.00 processors (2.00 for Partition A, 1.00 for Partition B, and 1.00 for Partition C). The managed system has only three IBM i licenses, but requires a total of four IBM i licenses, so the managed system is out of compliance with the IBM i license agreement.

If you have a managed system that can use multiple shared processor pools, you can use the Hardware Management Console (HMC) to configure a shared processor pool with a maximum processing unit value of 3.00, and assign Partition A, Partition B, and Partition C to that shared processor pool. If you do this, Partition A, Partition B, and Partition C can continue to be uncapped. You would remain in compliance with the IBM i license agreement because the maximum processing unit value would ensure that IBM i uses no more than three processing units.

For another example of a logical partition configuration that is out of compliance with a licensing agreement, the system administrator at Company Y changes the sharing mode of Partition B and Partition C back to capped. However, the system administrator then moves 0.50 processing units from Partition D to Partition A. Before the system administrator is allowed to do this, the system administrator must increase the number of virtual processors on Partition A from 2 to 3. The following table shows the new logical partition configuration.

Table 16. Logical partition configuration out of compliance with license agreement (second example)

Logical partition name	Operating system	Processing mode	Sharing mode	Processing units	Virtual processors	Maximum number of processors that can be used by the logical partition
Partition A	IBM i	Shared	Uncapped	2.25	3	3.00 (the number of virtual processors for the uncapped shared logical partition)
Partition B	IBM i	Shared	Capped	0.60	1	0.60 (the number of processing units for the capped shared logical partition)
Partition C	IBM i	Shared	Capped	0.40	1	0.40 (the number of processing units for the capped shared logical partition)
Partition D	Linux	Shared	Uncapped	0.75	2	2.00 (the number of virtual processors for the uncapped shared logical partition)

In this configuration, the three IBM i logical partitions can use a maximum of 4.00 processors (3.00 for Partition A, 0.60 for Partition B, and 0.40 for Partition C). The managed system has only three IBM i licenses, but requires a total of four IBM i licenses, so the managed system is out of compliance with the IBM i license agreement.

Considerations other than licensed program agreements might constrain your ability to run IBM licensed programs on certain server models.

Related concepts

Processors
A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

Minimum hardware configuration requirements for logical partitions

Each logical partition requires at least a certain amount of hardware resources. You can assign hardware resources directly to a logical partition, or you can set the logical partition to use the hardware resources that are assigned to another logical partition. The minimum hardware configuration requirements for each logical partition depend on the operating system or software that is installed on the logical partition.

The following table lists the minimum hardware requirements for logical partitions.

Table 17. Minimum hardware requirements for logical partitions

Minimum requirement	AIX and Linux	IBM i
Processor	One dedicated processor or 0.1 processing unit, or 0.05 processing unit when the firmware is at level FW760, or later.	One dedicated processor or 0.1 processing unit, or 0.05 processing unit when the firmware is at level FW760, or later. **Note:** HEA is not supported on POWER9 processor-based server.
Memory (physical or logical)	• AIX 5.3 to AIX 6.0: 128 MB • AIX 6.1 or later: 256 MB • Linux: 128 MB	2 GB plus 40 MB for each active logical Host Ethernet Adapter (LHEA)

Table 17. Minimum hardware requirements for logical partitions (continued)

Minimum requirement	**AIX and Linux**	**IBM i**
I/O	• Physical or virtual storage adapter (SCSI card) • Physical or virtual network adapter • Storage: – AIX: 2 GB – Linux: Approximately 1 GB	• Load source – Physical or virtual disk I/O adapter (IOA) – Physical or virtual disk unit that is at least 17 GB • Console: your choice of one of the following console types: – Hardware Management Console (HMC) 5250 emulation (requires an HMC) – Operations Console: Requires a LAN connection that supports Operations Console connections. The Operations Console LAN connection can be an embedded port or a LAN IOA. • Alternative restart device: your choice of tape or optical. These devices connect to either a physical or virtual storage adapter. The optical device can be a physical or virtual device. • Physical or virtual LAN adapter that can be used for serviceable event reporting and connection monitoring. At least one IBM i logical partition in the managed system must have a physical LAN adapter that the IBM i logical partition can use for serviceable event reporting and connection monitoring. You can then create a virtual LAN that connects the IBM i logical partition with the physical LAN adapter to the other logical partitions on the managed system, and bridge the physical LAN adapter to the virtual LAN. If the system is managed by using an HMC, the physical LAN adapter must be able to communicate with the HMC so that serviceable event reporting can be routed through the HMC.

Partitioning with the HMC

The *Hardware Management Console (HMC)* is a system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and forward information to IBM for analysis.

The HMC features a browser-based user interface. You can use the HMC locally by connecting a keyboard and mouse to the HMC. You can also configure the HMC so that you can connect to the HMC remotely using a supported browser.

Creating logical partitions

You can create an AIX, Linux, or IBM i logical partition by clicking **Create Partition**, or by using the **Create a Partition from Template** wizard.

About this task

For more information about creating a logical partition by using the **Create a Partition from Template** wizard, see Creating a logical partition by using a template.

To create an AIX, Linux, or IBM i logical partition by using the **Create Partition** option, see Creating logical partitions by using Create partition .

What to do next

When the Hardware Management Console (HMC) is at Version 9.1.0, or later, you can use the **mksyscfg** command to create a logical partition that supports the Physical Page Table (PPT) ratio. During Live Partition Mobility, the Physical Page Table (PPT) ratio is used to translate effective addresses to physical real addresses. PPT is the ratio of the physical memory of the partition and is used by the hypervisor for paging during Live Partition Mobility. To view the PPT ratio attribute of the logical partition, run the **lshwres** command.

Related information

Changing a partition template to disable Live Partition Mobility

Viewing system event logs for the Live Partition Mobility disable operation

Creating additional logical partitions

You can use the Create Logical Partition wizard on the Hardware Management Console (HMC) to create a new logical partition. When you create a logical partition, you also create a partition profile that contains the resource allocations and settings for the logical partition.

Before you begin

Use this procedure only if you are creating logical partitions on a managed system that has already been partitioned. If you are creating logical partitions on a new or nonpartitioned managed system, you must test the hardware on your managed system to ensure that the hardware is in working order. Testing the hardware helps you detect potential problems with your managed system and makes such problems easier to correct.

If you plan to create logical partitions that use shared memory, you must first configure the shared memory pool. For instructions, see "Configuring the shared memory pool" on page 113.

If you plan to create AIX logical partitions that use Active Memory Expansion, you must first enable Active Memory Expansion for the server by entering an activation code. For instructions, see "Entering the activation code for Active Memory Expansion" on page 117.

If you want to assign single root I/O virtualization (SR-IOV) logical ports to a logical partition during partition creation, verify whether the managed system supports SR-IOV before you create the logical partition.

About this task

For more information about creating a logical partition and a partition profile on your server by using the HMC, see "Creating logical partitions" on page 79

What to do next

After creating your logical partition and partition profile, you must install an operating system. For installation instructions for the AIX, IBM i, and Linux operating systems, see Working with operating systems and software applications for POWER9 processor-based systems. For installation instructions for the Virtual I/O Server, see Installing the Virtual I/O Server and client logical partitions.

Related information

Changing a partition template to disable Live Partition Mobility

Viewing system event logs for the Live Partition Mobility disable operation

Creating logical partitions on a new or nonpartitioned server

Use these procedures to create logical partitions on your new or nonpartitioned server using the Hardware Management Console (HMC).

When you receive your server, the server is in what is known as the manufacturing default configuration. You can install an operating system on the server and use the server in a nonpartitioned configuration. However, if you want to create logical partitions on the managed system, you must develop a logical partition plan for the server, add hardware to the server or move the hardware within the server according to your logical partition plan, and validate the hardware on the server. When the server is ready, you can then create the logical partitions using the HMC.

The procedure used to create logical partitions on a new or nonpartitioned server varies by server type.

Assigning a single root I/O virtualization logical port to a logical partition

You can assign a single root I/O virtualization (SR-IOV) logical port to a logical partition by using the Hardware Management Console (HMC).

About this task

For information about managing hardware virtualized I/O adapters, see Managing hardware virtualized I/O adapters.

Creating a logical partition with synchronization of the current configuration

You can create an AIX or LinuxLinux logical partition with synchronization of the current configuration capability by using the Hardware Management Console (HMC).

Prerequisites and assumptions

Ensure that the following prerequisite tasks have been completed before you start the configuration steps:

1. The HMC is set up and configured. For instructions, see Installing and configuring the HMC.
2. You read and understand the “Logical partition overview” on page 2.
3. You completed the tasks recommended for logical partition planning. For instructions, see “Planning for logical partitions” on page 61.
4. You removed the system from the manufacturing default configuration and moved the physical hardware to support a partitioned configuration.
5. You have logged in to the HMC with one of the following user roles:
   - Super administrator
   - Operator
6. Ensure that the HMC is at Version 7 Release 7.8.0, or later.

Configuration steps

For more information about creating a logical partition when the HMC is at version 8.7.0, or later, see “Creating logical partitions” on page 79.

When you create a logical partition by using the **Create Partition** option, or by using the **Create a Partition from Template** wizard, the synchronization of the current configuration is enabled by default.

Enabling the synchronization of the current configuration capability

You can enable the synchronization of the current configuration capability on a logical partition by using the Hardware Management Console (HMC), after the logical partition is created.

Before you begin

Before you plan to enable the feature, ensure that the HMC is at Version 7 Release 7.8.0, is later.

About this task

For information about enabling the synchronization of the current configuration capability on a logical partition after the logical partition is created, see Changing advanced partition settings.

Remote restart states

A simplified remote restartable partition goes through several state changes regarding the simplified remote restart operation, both on the source and destination servers. Most of the simplified remote restart operations are supported only when the partition is in the appropriate remote restart state. A remote restart state is not related to the logical partition state, but is an indicator that is specifically associated with the simplified remote restart operation.

You can use the **`lssyscfg`** command to view the remote restart status of the partition. The following are the possible values:

Invalid
The partition that is configured for simplified remote restart is in the `Invalid` state until the logical partition is activated. Simplified remote restart is supported only on a logical partition that has been started at least once.

Remote Restartable
After the partition is started and running, the partition transitions into the `Remote Restartable` state. A partition in this state can be remote restarted.

Source Remote Restarting
During the actual simplified remote restart operation, the source partition is in the `Source Remote Restarting` state. This state is transitional and valid until the simplified remote restart operation completes or the operation is canceled.

Destination Remote Restarting
During the actual simplified remote restart operation, the destination partition transitions into the `Destination Remote Restarting` state. This state is transitional and is valid until the simplified remote restart operation completes or the operation is canceled.

Destination Remote Restarted
When the simplified remote restart operation reaches the no return point on the destination system (all the adapters are configured on the destination system), the remote restart status is set to `Destination Remote Restarted`.

Remote Restarted
After the simplified remote restart operation completes, the source logical partition is in the `Remote Restarted` state. The source logical partition can be cleaned up and the destination logical partition is now again ready to be restarted as needed.

Local Storage Update Failed
When an update to persisted information (configuration data stored external to the server on persistent storage) on the Hardware Management Console (HMC) fails due to any reason, the logical partition is in the `Local Storage Update Failed` state. This state indicates that the persisted information on the HMC is out of synchronization with the current logical partition configuration. Simplified remote restart is not allowed in this remote restart state. However, you can use the *usecurrdata* option with the **`rrstartlpar`** command to run a simplified remote restart operation.

Forced Source Side Restart
When you use the *usecurrdata* option with the **`rrstartlpar`** command to run a simplified remote restart operation, the partition is restarted with the configuration data on the device and the remote restart state on the source system is updated to `Forced Source Side Restart` state.

Partial Update
When a system is connected to an HMC that has logical partitions with the simplified remote restart capability enabled, the HMC automatically collects the configuration information and the data is stored external to the server on persistent storage. Some configuration information like virtual adapter information, requires a Resource Monitoring and Control (RMC) connection to the Virtual I/O Server (VIOS) partitions. Therefore, the HMC waits until the RMC connection is established to collect such information. When the virtual adapter information is not collected for any reason, the remote restart state is set to `Partial Update`.

Stale Data

When a system is connected to an HMC, the remote restart state is set to `Stale Data` if there is configuration information existing for a logical partition before the state is changed to `Partial Update`.

Out Of Space

When an update to the persisted information fails because there is insufficient space on storage disk of the HMC to store the configuration information, the remote restart state is updated to `Out Of Space`. You can free up space on storage disk of the HMC and run the **`refdev`** command to recover from this state.

Profile Restored

When a profile restore operation is performed on a system, during the creation of the simplified remote restart capable partition, the remote restart state is set to `Profile Restored`.

Source Side Cleanup Failed

When a cleanup operation performed on the source system after a successful simplified remote restart fails, remote restart state on the source partition is set to `Source Side Cleanup Failed`.

Reserved Storage Device Update Failed

This state is specific to the simplified remote restart operation that requires a reserved storage device. When an update to the reserved storage device fails for any reason, the logical partition is in the `Reserved Storage Device Update Failed` state. This state indicates that the data on the device is not synchronized with the current partition configuration. Simplified remote restart is not allowed in this remote restart state. However, you can use the *-force* option to run a simplified remote restart operation.

Recovering a simplified remote restart operation

If a simplified remote restart operation of an AIX, Linux, or IBM i logical partition fails, the Hardware Management Console (HMC) attempts an auto recover operation. When the auto recover operation fails, you can recover a simplified remote restart operation by using the HMC command-line interface.

Procedure

To recover a simplified remote restart operation, on the HMC command line, type the following command:

```
rrstartlpar -o recover -m <source server> -t <destination server> -p <lpar name>
| --id <lpar id> [--force]
```

Aborting a simplified remote restart operation

You can abort or cancel a simplified remote restart operation of an AIX, Linux, or IBM i logical partition by using the Hardware Management Console (HMC) command-line interface.

Procedure

To abort or cancel the simplified remote restart operation of a logical partition, type the following command from the HMC command line:

```
 rrstartlpar -m <source managed system> -t <target manged system> -p <lpar name>
 -ip <IP address> [-u <user ID>] -o cancel
```

What to do next

After the cancel operation completes on the destination server, the remote restart status on source server changes to **Remote Restartable**.

Related information

rrstartlpar command

Viewing the details of a simplified remote restart operation

You can view the details of a simplified remote restart operation of an AIX, Linux, or IBM i logical partition, or the server by using the Hardware Management Console (HMC) command-line interface.

Before you begin

You must ensure that the server is in the `Standby` or `Operating` states when you are running the **`lsrrstartlpar`** command for server-level details. The **`lsrrstartlpar`** command can be run at the logical partition level for all system states that support the simplified remote restart feature.

Procedure

1. To view the details of the simplified remote restart operation of the server, type the following command from the HMC command line:

   ```
   lsrrstartlpar -r sys -m <managed system>
   ```

2. To view the details of the simplified remote restart operation of a logical partition, type the following command from the HMC command line:

   ```
   lsrrstartlpar -r lpar -m <managed system> [--filter "lpar_names=" | "lpar_ids=""""]
   ```

3. To verify whether the auto cleanup feature is enabled, type the following command from the HMC command line:

   ```
   lsrrstartlpar -r -mc <managed system>
   ```

4. To view the configuration information of all logical partitions that support the simplified remote restart feature, type the following command from the HMC command line:

   ```
   lsrrstartlpar -r lparcfg -m <managed system>
   ```

   Where *lparcfg* shows the logical partition configuration details from the simplified remote restart data.

5. You can view the possible and the recommended mappings of the virtual Fibre Channel adapters, virtual SCSI adapters, and virtual Network Interface Controllers (vNICs) of a logical partition, to Virtual I/O Servers on the destination server, by typing the following command from the HMC command line:

   ```
   lsrrstartlpar -r virtualio -o validate  -m <source-system-name> -t <target-system-name>
   [--vniccfg <1|2> ] [-- redundantvnicbkdev <1|2>]
   [--ip <IP address>]  [-u <user ID> ]
    --filter "<filter data>"
   ```

   Where

   - *virtualio* lists all the available Virtual I/O adapters on the destination system that can be used for mapping virtual I/O.
   - *vniccfg* specifies whether the vNIC configuration must be maintained. A value of 1 indicates that the vNIC configuration must be maintained. If you specify the value as 1, and if the vNIC configuration cannot be maintained, the remote restart operation fails. A value of 2 indicates that the vNIC configuration must be maintained when possible. If you specify the value as 2, and if the vNIC configuration cannot be maintained, the remote restart operation succeeds but the vNIC configuration of the logical partition is not retained after the remote restart operation completes.
   - *redundantvnicbkdev* specifies whether the vNIC backing device redundancy must be maintained. A value of 1 indicates that the vNIC backing device redundancy must be maintained. If you specify the value as 1, and if the vNIC backing device redundancy cannot be maintained, the remote restart operation fails. A value of 2 indicates that the vNIC backing device redundancy must be maintained when possible. If you specify the value as 2, and if the vNIC backing device redundancy cannot be maintained, the remote restart operation succeeds but the vNIC backing device redundancy of the logical partition is not retained after the remote restart operation completes.
   - *ip* is the IP address or the host name of the management console of the destination server.

- *u* is the user ID that must be used on the management console of the destination server.
- *filter* is a required parameter and only one logical partition (either partition ID or partition name) can be specified when the *virtualio* parameter is specified.

Related information

lsrrstartlpar command

Validating the simplified remote restart operation of a logical partition

You can validate the simplified remote restart operation of an AIX, Linux, or IBM i logical partition by using the Hardware Management Console (HMC) command-line interface.

Before you begin

When the source and destination servers are managed by different Hardware Management Consoles, you must ensure that the HMC at the source server and the destination server is at version 8.5.0 or later.

Note: When the HMC is at Version 9.2.950, or later, and the firmware is at a level FW950, or later, for the remote restart operation to be successful on the logical partition that supports platform keystore capability, the user-defined system key that is configured on both the source and the destination system must match.

Procedure

1. To validate the simplified remote restart operation of a logical partition, type the following command from the HMC command line:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> |
   --id <lpar id> -o validate
   ```

   To validate the simplified remote restart operation of a logical partition when the source server and the destination server are managed by a different HMC, type the following command:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> |
   --id <lpar id>  --ip <IP address> [-u <user ID>] -o validate
   ```

   Where

   - *IP address* is the IP address or the host name of the HMC that manages the destination server.
   - *user ID* is the user ID that is used on the HMC that manages the destination server.

   If the HMC at the destination server is not at version 8.5.0 or later, the validation operation fails and error messages are displayed for your appropriate action.
2. To specify the destination shared processor pool, type the following command from the HMC command line:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>
   | --id <lpar id> --ip <IP address> [-u <user ID>] - i|f
   “shared_proc_pool_name=<spp name>|shared_proc_pool_id=<spp id>” -o validate
   ```

   Where

   - *shared_proc_pool_name* is the name of the shared processor pool on the destination server.
   - *shared_proc_pool_id* is the ID of the shared processor pool.

   Note: The *shared_proc_pool_id* and *shared_proc_pool_name* attributes are mutually exclusive.
3. To specify the destination virtual Fibre Channel mapping, type the following command from the HMC command line:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>
   | --id <lpar id> --ip  <IP address> [-u <user ID>] - i|f  “virtual_fc_mappings
   =slot_num/vios_lpar_name/vios_lpar_id/[vios_slot_num]/[vios_fc_port_name]”  -o validate
   ```

   Where

   - *slot num* is virtual Fibre Channel slot number.
   - *vios_slot_num* is virtual Fibre Channel slot number of the VIOS.

4. To validate the simplified remote restart operation of a logical partition with a different number of dedicated processors or virtual processors on the destination server than what the partition was assigned on the source server, or with a different number of shared processing units on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

   ```
   rrstartlpar -o validate -m <managed system> -t <target system> [-p <partition name>
   | --id <partition ID>] -i “desired_procs = <desired procs absolute value>|min,
   desired_proc_units = <desired proc units absolute value>|min”
   ```

   The value for the *desired_procs* attribute indicates the number of processors with which logical partition can be restarted. The following are the possible values for this attribute:

   - *desired procs absolute value* - Must be greater than or equal to the current minimum dedicated processors or virtual processors, and less than or equal to current maximum dedicated processors or virtual processors.
   - *desired proc units absolute value* - Must be greater than or equal to the current minimum processing units, and less than or equal to current maximum processing units.
   - *min* - The value of the current minimum processing units, or the current minimum dedicated processors or virtual processors. In the case of dedicated processors, the logical partition is started with the minimum number of dedicated processors or virtual processors the logical partition was assigned on the source server. In the case of shared processors, the logical partition is started with the minimum number of shared processing units the partition was assigned on the source server.

5. To validate the simplified remote restart operation of a logical partition when you want to start the logical partition with a different amount of memory on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

   ```
   rrstartlpar -o validate -m <managed system> -t <target system> [-p <partition name>
    | --id <partition ID>]  -i “desired_mem = <desired absolute memory>|min”
   ```

   ```
   
   ```

   The value for the *desired_mem* attribute indicates the amount of memory with which the logical partition can be restarted. The following are the possible values for this attribute:

   - *desired_absolute_memory* - Must be greater than or equal to the current minimum memory, and less than or equal to current maximum memory.
   - *min* - The value of the current minimum memory. The logical partition is started with the minimum amount of memory the logical partition was assigned on the source server.

6. To validate the simplified remote restart operation of a logical partition when you want to start the logical partition with a different virtual switch on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

   ```
   rrstartlpar -o validate -m <managed system> -t <target system> [-p <partition name>
   | --id <partition ID>] -i “vswitch_mappings = <vlanId_1>/<source_vswitchName_1>/
   <target_vswitchName_1>, …, <vlanId_n>/<source_vswitchName_n>/<target_vswitchName_n>”]”
   ```

Related information

rrstartlpar command

Remotely restarting a logical partition

You can restart an AIX, Linux, or IBM i logical partition remotely by using the Hardware Management Console (HMC) command-line interface. You can run up to four concurrent simplified remote restart operations for a destination server. When the HMC is at version 8.5.0 or later, you can run up to 32 concurrent simplified remote restart operations for a destination server.

Before you begin

Before you begin, complete the following tasks:

- You must ensure that the HMC is at version 8.2.0 or later, and the firmware is at level FW820, or later, and the VIOS must be at version 2.2.3.4 with the VIOS interim fix IV63331m4a or later, for the simplified remote restart feature.
- You must ensure that the source and the destination servers are connected to the same HMC. However, when the HMC is at version 8.5.0 or later, the source and destination server can be connected to a different Hardware Management Consoles. In this case, the HMC at the source and destination servers must be at version 8.5.0, or later.
- When the HMC is at version 8.3.0, or later, you can restart a logical partition on another server only when the source server is in the `Initializing`, `Power Off`, `Error`, or `Error - Dump in progress` state and the destination server is in the `Operating` state. When the HMC is at version 8.4.0, or later, you can restart a logical partition on another server only when the source server is in the `Initializing`, `Power Off`, `Powering Off`, `Error`, or `Error - Dump in progress` in progress state and the destination server is in the `Operating` state. When the HMC is at version 8.5.0, or later, you can restart a logical partition on another server only when the source server is in the `Initializing`, `Power Off`, `Powering Off`, `No Connection`, `Error`, or `Error - Dump in progress` state, and the destination server is in the `Operating` state.
- When the HMC is at Version 9.2.950, or later, and the firmware is at a level FW950, or later, for the remote restart operation to be successful on the logical partition that supports platform keystore capability, the user-defined system key that is configured on both the source and the destination system must match.

Procedure

1. To remotely restart a logical partition, type the following command from the HMC command line:

   ```
    rrstartlpar -o restart -m <source managed system> -t <target manged system> -p <lpar name>
   ```

   When the source server and the destination server are managed by different Hardware Management Consoles, type the following command to remotely restart a logical partition:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> | --id
   <lpar id>
   --ip <IP address> [-u <user ID>] -o restart
   ```

   Note: If a logical partition that is enabled with the simplified version of the remote restart feature was in the Suspended state before the simplified remote restart operation started, the simplified remote restart operation fails. You can use the *--force* option to force a simplified remote restart operation.

   To verify the status of the simplified remote restart operation, type the following command from the HMC command line:

   ```
   lssyscfg -r lpar -m <server> -F name,state,remote_restart_status
   ```

   You can also run the **`lsrrstartlpar`** command to view the status of the simplified remote restart operation.
2. To specify the destination shared processor pool for the simplified remote restart operation, type the following command from the HMC command line:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>
   | --id <lpar id> --ip <IP address> [-u <user ID>] - i|f
   “shared_proc_pool_name=<spp name>|shared_proc_pool_id=<spp id>” -o restart
   ```

   Where:

   - *shared_proc_pool_name* is the name of the shared processor pool on the destination server.
   - *shared_proc_pool_id* is the ID of the shared processor pool.

Note: The *shared_proc_pool_id* and *shared_proc_pool_name* attributes are mutually exclusive.

3. To specify the destination virtual Fibre Channel mapping for the simplified remote restart operation, type the following command from the HMC command line:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> | --id
   <lpar id>
   --ip <IP address> [-u <user ID>] - i|f “virtual_fc_mappings=
   slot_num/vios_lpar_name/vios_lpar_id/[vios_slot_num]/[vios_fc_port_name]” -o restart
   ```

   Where:

   - *slot num* is virtual Fibre Channel slot number.
   - *vios_slot_num* is virtual Fibre Channel slot number of the VIOS.

4. To remotely restart a partition when the source server is in the `No Connection` state, type the following command from the HMC command line:

   ```
   rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> | --id
   <lpar id> -o restart --noconnection
   ```

5. When the HMC is at version 8.5.0, or later, an auto cleanup operation on the server is performed by the HMC after the successful completion of the simplified remote restart operation. The source server is back to the `Standby` or `Operating` state, and the Resource Monitoring and Control (RMC) connection for the Virtual I/O Server (VIOS) partitions that are serving the remotely restarted logical partitions becomes active.

   You can enable the auto cleanup operation by typing the following command from the HMC command line:

   ```
   rrstartlpar -o set -r <source managed system> -i "auto_cleanup_enabled=1"
   ```

   You can disable the auto cleanup operation by typing the following command from the HMC command line:

   ```
   rrstartlpar -o set -r <source managed system> -i "auto_cleanup_enabled=0"
   ```

   The auto cleanup operation is enabled by default. If the remote restarted partitions are not cleaned up automatically, type the following command from the HMC command line to perform a cleanup operation on the source server:

   ```
   rrstartlpar -o cleanup -m source managed system -p lpar name
   ```

   The default behavior of a cleanup operation is to remove the reserved storage device from the device pool on the source server. You can use the *--retaindev* option when you want to retain the reserved storage device in the pool and override the default action of the cleanup operation.

6. To remotely restart a logical partition with a different number of dedicated processors or virtual processors on the destination server than what the partition was assigned on the source server, or with a different number of shared processing units on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

   ```
   rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name>
   | --id <partition ID>] -i “desired_procs = <desired procs absolute value>|min,
   desired_proc_units =
   <desired proc units absolute value>|min”
   ```

   The value for the *desired_procs* attribute indicates the number of processors with which the logical partition can be restarted. The following are the possible values for this attribute:

   - *desired procs absolute value* - Must be greater than or equal to the current minimum dedicated processors or virtual processors, and less than or equal to current maximum dedicated processors or virtual processors.
   - *desired proc units absolute value* - Must be greater than or equal to the current minimum processing units, and less than or equal to current maximum processing units.

- *min* - The value of the current minimum processing units, or the current minimum dedicated processors or virtual processors. In the case of dedicated processors, the logical partition is started with the minimum number of dedicated processors or virtual processors the logical partition was assigned on the source server. In the case of shared processors, the logical partition is started with the minimum number of shared processing units the partition was assigned on the source server.

7. To remotely restart a logical partition when you want to start the logical partition with a different amount of memory on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

   ```
   rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name> |
    --id <partition ID>] -i “desired_mem = <desired absolute memory>|min”
   ```

   The value for the *desired_mem* attribute indicates the amount of memory with which the logical partition can be restarted. The following are the possible values for this attribute:

   - *desired_absolute_memory* - Must be greater than or equal to the current minimum memory, and less than or equal to current maximum memory.
   - *min* - The value of the current minimum memory. The logical partition is started with the minimum amount of memory the logical partition was assigned on the source server.

8. To remotely restart a logical partition when you want to start the logical partition with a different virtual switch on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

   ```
   rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name> |
   --id <partition ID>] -i “vswitch_mappings = <vlanId_1>/<source_vswitchName_1>/
   <target_vswitchName_1>, …, <vlanId_n>/<source_vswitchName_n>/<target_vswitchName_n>”]”
   ```

9. When you want to prevent a logical partition from being started during the simplified remote restart operation, type the following command from the HMC command line:

   ```
   rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name>
    | --id <partition ID>] --skippoweron
   ```

10. You can test the simplified remote restart operation of a shutdown partition from the source server, which is in `Operating` or `Standby` state by typing the following command from the HMC command line:

    ```
    rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name>
    | --id <partition ID>] --test
    ```

Related information

rrstartlpar command

Enabling or disabling the simplified remote restart capability

You can enable or disable the simplified remote restart capability of a logical partition after the logical partition is created, by using the Hardware Management Console (HMC).

Before you begin

- Ensure that the HMC is at Version 8.1.0, or later.
- Ensure that the server is a POWER9 processor-based server that supports partitions that are capable of simplified remote restart.
- Ensure that the logical partition is in the `Not Activated` state. When the HMC is at Version 8.6.0, or later, and the firmware is at level FW860, or later, you can enable or disable the simplified remote restart capability when the logical partition is in the `Running` state. The logical partition must not be in the `Suspended`, `Resuming`, `Migrating`, or `Remote Restarting` states.
- To use the simplified remote restart feature, ensure that the HMC is at Version 8.2.0, or later.

About this task

To enable or disable the simplified remote restart capability of a logical partition by using the HMC, complete the following steps:

Procedure

1. From the HMC command line, type the following command to enable the simplified remote restart feature:

   ```
    chsyscfg -r lpar -m managed-system -i "name=partition name, remote_restart_capable=1"
   ```

   From the HMC command line, type the following command to enable the simplified remote restart feature:

   ```
    chsyscfg -r lpar -m managed-system -i "name=partition name,
   simplified_remote_restart_capable=1"
   ```

2. From the HMC command line, type the following command to disable the simplified remote restart feature:

   ```
    chsyscfg -r lpar -m managed-system -i "name=partition name, remote_restart_capable=0"
   ```

   From the HMC command line, type the following command to disable the simplified remote restart feature:

   ```
    chsyscfg -r lpar -m managed-system -i "name=partition name,
   simplified_remote_restart_capable=0"
   ```

Creating a logical partition with the simplified remote restart capability

You can create an AIX, IBM i, or Linux logical partition with the simplified remote restart capability by using the Hardware Management Console (HMC). The HMC provides options to enable the simplified remote restart of the logical partition when the logical partition is created.

Prerequisites and assumptions

Ensure that the following prerequisite steps have been completed before you start the configuration steps:

- Before you create a logical partition with the simplified remote restart capability, verify the following requirements:
  1. You read the configuration requirements and restrictions for creating a logical partition with the simplified remote restart capability. For instructions, see “Configuration requirements and restrictions for the remote restart capability of a logical partition ” on page 67.
  2. Before you create a logical partition with the simplified remote restart capability, verify the following requirements:
     - Verify that the server supports the simplified remote restart capability. For instructions, see “Verifying that the server supports partitions that are capable of the simplified version of the remote restart feature” on page 68.
     - You read the configuration requirements and restrictions for creating a logical partition with the simplified remote restart capability. For instructions, see “Configuration requirements and restrictions for the remote restart capability of a logical partition ” on page 67.

You can create a partition with the simplified remote restart capability by using a template that supports the capability when you create a partition by using the **Create a Partition from Template** wizard. When the template does not support the capability, you can enable the capability by modifying the template and then using the modified template to create a partition by using the **Create a Partition from Template** wizard. For more information about editing a partition template, see Changing a partition template.

For more information about enabling the simplified remote restart feature after you create the logical partition, see Changing partition properties and capabilities.

Creating a logical partition with Virtual Trusted Platform capability

You can create an AIX logical partition with Virtual Trusted Platform Module (VTPM) capability by using the Hardware Management Console (HMC). HMC Version 7.7.4 or later, provides an option to enable a VTPM on the logical partition when the logical partition is created. The HMC also provides an option to enable a VTPM on a running logical partition.

Prerequisites and assumptions

Ensure that the following prerequisite tasks have been completed and are operational before you start the configuration steps:

1. The HMC is set up and configured. For instructions, see Installing and configuring the HMC.
2. You read and understand the “Logical partition overview” on page 2.
3. You completed the tasks recommended for logical partition planning. For instructions, see “Planning for logical partitions” on page 61.
4. You removed the system from the manufacturing default configuration and moved the physical hardware to support a partitioned configuration. For instructions, see “Creating logical partitions on a new or nonpartitioned server” on page 80.
5. You verified that the server has logical partition support for VTPM. For instructions, see “Verifying that the server supports Virtual Trusted Platform Module ” on page 69

   You can also enable the VTPM on a logical partition after logical partition creation. For instructions, see “Enabling and disabling a Virtual Trusted Platform Module on a logical partition ” on page 91
6. You have logged in to the HMC with one of the following user roles:
   - Super administrator
   - Operator

You can create a partition with the VTPM capability by using a template that supports the capability when you create a partition by using the **Create a Partition from Template** wizard. When the template does not support the capability, you can enable the capability by modifying the template and then using the modified template to create a partition by using the **Create a Partition from Template** wizard. For more information about editing a partition template, see Changing a partition template. For more information about enabling the VTPM capability after partition creation, see Changing partition properties and capabilities.

Enabling and disabling a Virtual Trusted Platform Module on a logical partition

You can enable a Virtual Trusted Platform Module (VTPM) on a logical partition by using the Hardware Management Console (HMC), after the logical partition is created.

Before you begin

To enable a VTPM, ensure that an AIX, Linux or a Virtual I/O Server (VIOS) logical partition is in the `Not activated` state.

About this task

For information about enabling VTPM on a logical partition, see Changing advanced partition settings.

Results

If you dynamically enable a VTPM on a logical partition, the VTPM function is activated only at the next logical partition activation. However, disabling a VTPM takes effect immediately.

What to do next

To dynamically disable a VTPM, log on to the AIX, Linux or VIOS logical partition and disable the Trusted Computing Services daemon (tcsd) by using the **`stopsrc`** command. When the **`tcsd`** software is stopped, the device must be removed from the AIX logical partition by using the **`rmdev`** command. After the device is successfully deleted from the AIX logical partition, use the HMC to clear the VTPM check box from the properties of the partition. This completely removes the device and deletes all stored data that is associated with the VTPM.

Viewing the Virtual Trusted Platform Module settings

You can view the advanced Virtual Trusted Platform Module (VTPM) settings by using the Hardware Management Console (HMC).

About this task

For more information about viewing the VTPM settings, see Changing advanced partition settings.

Creating additional partition profiles

You can create more than one partition profile for a logical partition using the Hardware Management Console (HMC). Each partition profile can specify a different amount of system resources and different logical partition startup attributes. You can change the attributes used by a logical partition by shutting down the logical partition and restarting the logical partition using a different partition profile.

Before you begin

If you plan to create a partition profile in which you configure Active Memory Expansion for an AIX logical partition, ensure that you enter an activation code to enable Active Memory Expansion on the server before you activate the logical partition with this partition profile. For instructions, see "Entering the activation code for Active Memory Expansion" on page 117.

When you create a partition profile, do not select **Use all the resources in the system** when both of the following conditions are true:

- You plan to create a partition profile that uses all of the system resources.
- You plan to configure Active Memory Expansion for that partition profile.

Instead, manually assign all of the resources in the system to the partition profile. In the process, you can configure the Active Memory Expansion factor.

About this task

To create a partition profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Click **Actions** > **New**.
5. Follow the steps in the Create Partition Profile wizard to create the partition profile.

What to do next

If you created at least one virtual fibre channel adapter, complete the following tasks to connect the logical partition to its storage:

1. Activate the logical partition. When you activate the logical partition, the HMC assigns a pair of worldwide port names (WWPNs) to the virtual fibre channel adapter. For instructions, see "Activating a logical partition" on page 126.
2. Restart the Virtual I/O Server (that provides the connection to a physical fibre channel adapter) or run the **syscfg** command. This enables the Virtual I/O Server to recognize the WWPNs of the virtual fibre channel adapter on the client logical partition For instructions, see "Restarting Virtual I/O Server logical partitions by using the HMC" on page 138.
3. Assign the virtual fibre channel adapter on the client logical partition to a physical port of a physical fibre channel adapter. For instructions, see "Changing virtual Fibre Channel for a Virtual I/O Server by using the HMC" on page 166.

Creating a system profile

You can create a system profile using the Hardware Management Console (HMC). A *system profile* is an ordered list of partition profiles. When you activate a system profile, the managed system attempts to activate the partition profiles in the system profile in the order in which the partition profiles are listed.

Before you begin

System profiles are also useful for validating your partition profiles to ensure that you have not overcommitted the resources on your managed system.

Restriction: You cannot create system profiles that contain logical partitions that use shared memory.

About this task

To create a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
5. Click **Actions** > **New**.
6. Enter the name of the new system profile into **System profile name**.
7. For each partition profile that you want to add to the system profile, open the logical partition to which the partition profile belongs, select the partition profile, and click **Add**.
8. Click **OK**.

Creating an AIX logical partition that uses IBM i virtual I/O resources

You can create an AIX logical partition that uses IBM i virtual I/O resources on servers that are managed by a Hardware Management Console (HMC). This allows you to maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Before you begin

To set this up, you must create virtual SCSI adapters that connect the AIX logical partition with the IBM i. You can then set up IBM i to provide disk resources to the AIX logical partition through the virtual SCSI connection. You can also create a virtual serial connection between the IBM i logical partition and the AIX logical partition. A virtual serial connection allows you to connect to the AIX logical partition from the IBM i logical partition.

Alternatively, you can create a Virtual I/O Server logical partition and configure the AIX logical partition to use the virtual SCSI and virtual Ethernet resources of the Virtual I/O Server logical partition. You might need to enter a PowerVM EditionsPowerVM for IBM PowerLinux activation code to create a Virtual I/O Server logical partition on your server.

About this task

For more information about creating a logical partition, see "Creating logical partitions" on page 79. For more information about adding IBM i hosted virtual SCSI adapters when the HMC is at version 8.7.0, or later, see Managing virtual storage for a partition.

Creating a network-server description and a network-server storage space for an AIX logical partition

A *network-server description (NWSD)* is an IBM i object that describes the storage resources that are used by an integrated operating environment. An NWSD can be linked to one or more network-server storage spaces. Create an NWSD to assign storage to an AIX logical partition that uses IBM i resources.

About this task

To create an NWSD and a network-server storage space for an AIX logical partition that uses IBM i resources, follow these steps:

Procedure

1. Determine the correct SCSI server resource name.
   - If there is only one SCSI server adapter corresponding to a given client logical partition, and that adapter has its remote logical partition and remote slot configured correctly, you can specify `*AUTO` as the RSRCNAME in your NWSD.
   - Otherwise, you must determine the actual resource name. At an IBM i command line, type `WRKHDWRSC *CMN`, and find a controller resource with type 290B and a converged location code that corresponds to the SCSI server adapter at the Hardware Management Console (HMC). This resource name will be used later to specify the SCSI server resource.
2. At an IBM i command line on the logical partition that shares resources, type `CRTNWSD` and press F4 for prompts.
3. Specify the following information:

   The default or suggested parameter values are provided within the parentheses. These settings are relevant only to a logical partition. After the installation, if your root file system (/) is not installed on the first partition of the first disk, you must set a root parameter.

   - `NWSD (Provide a name for the NWSD)`
   - `RSRCNAME (*AUTO or the resource name of the SCSI server resource`
   - `TYPE(*GUEST)`
   - `ONLINE (*NO or *YES)`
   - `PARTITION ('Provide the name of your AIX logical partition')`

     As an alternative to the Partition parameter, you can also specify a logical partition number by typing `PTNNBR(`*`integer`*`)` where *integer* is the number of the logical partition you are specifying.
   - `CODEPAGE (437)`
   - `TCPPORTCFG (*NONE)`
   - `RSTDDEVRSC (for virtual CD and tape devices) (*NONE)`
   - `SYNCTIME (*TYPE)`
   - `IPLSRC (*NWSSTG)`

  - You can store a kernel in a disk partition of a virtual disk (a network-server storage space (NWSSTG)). By specifying the `IPLSRC (*NWSSTG)` parameter, you are specifying that the AIX logical partition will start from a disk partition on that virtual disk. The disk partition on the virtual disk must be formatted as type `PReP Boot` (type 0x41) and marked as a device that starts.
  - To start an NWSD with a kernel from a stream file, set the `IPLSRC` parameter to `*STMF` and set the `IPLSTMF` parameter to point to the kernel. You must have read access to the file and the path leading to the file to use the vary on command. This value only loads the kernel. After the kernel is running, it must find a root file system. In an initial installation, the root file system might be a RAM disk that is physically attached to the kernel.
- `IPLSTMF (*NONE)`
- `IPLPARM (*NONE)`
- `PWRCTL (*YES)`
  - If you specify `PWRCTL (*YES)`, perform the following steps:
    a. Ensure that the server adapter in the IBM i logical partition specifies the remote logical partition and remote slot in its configuration.
    b. Ensure that the client logical partition has the IBM i logical partition as the power-controlling logical partition in the profile.
    c. Ensure before you activate the NWSD that the client logical partition's profile has been saved to the server by activating the logical partition from the HMC, even if the client operating system does not activate correctly because of the absence of virtual devices.
  - If you specify `PWRCTL(*NO)`, virtual devices will be available to the logical partition. You must shut down and restart the logical partition using the HMC.

4. If you use IBM Navigator for i, create the network-server storage space using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration**.
   b) Right-click the **Disk Drives** and select **New Disk**.
   c) In the **Disk drive name** field, specify the name that you want to give to the disk drive.
   d) In the **Description** field, specify a meaningful description for the disk drive.
   e) In the **Capacity** field, specify the size of the new disk drive in megabytes.
   f) Click **OK**.
   g) Continue with step "6" on page 95.
5. If you use a character-based interface, create the network-server storage space using a character-based interface:
   a) At an IBM i command line, type the command CRTNWSSTG and press F4.
   The Create NWS Storage Space (CRTNWSSTG) display opens.
   b) In the Network-server storage space field, specify the name you want to give to the storage space.
   c) In the Size field, specify the size in megabytes for the new storage space.
   d) In the Text description field, specify a meaningful description for the storage space.
   e) Press Enter.
   f) Continue with step "7" on page 96.
6. If you use IBM Navigator for i, link the network-server storage space using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration**.
   b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
   c) Select the server to which you want to link the network-server storage space.
   d) Select the link sequence position you want to use.
   e) Select one of the available data access types.
   f) Click **OK**.

   The procedure is complete. Do not complete step "7" on page 96.

7. If you use a character-based interface, link the network-server storage space using a character-based interface:
   a) At an IBM i command line, type the command ADDNWSSTGL and press F4.

      The Add Network-Server Storage Link (ADDNWSSTGL) display opens.
   b) In the Network server description field, specify the name of the network server description (NWSD).
   c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the AIX logical partition).
   d) In the Drive sequence number field, specify the link sequence position you want to use.
   e) Press Enter.

Connecting an AIX logical partition to the virtual console

You can connect an AIX logical partition to the virtual console so that you can install the operating system or access the command line interface for the AIX logical partition.

Before you begin

You must have one of the following privileges to use the AIX virtual console:

- Remote Panel
- System Partitions - Administration

About this task

The virtual console provides the console function for an AIX server. It is used primarily during the initial installation of the operating system. The virtual console can also be used to view server errors or to restore communication to the LAN. This console connection is used prior to configuring TCP/IP.

Any Telnet client can be used as the AIX console. Multiple Telnet clients can share access to the same virtual console. To connect to a console, use Telnet to connect to port 2301 of the logical partition that is sharing its resources. TCP/IP must be configured and running on at least one IBM i logical partition. Complete one of the following procedures:

- If you use IBM Personal Communications, connect to a virtual console using IBM Personal Communications.
  1. Click **Start** > **IBM Personal Communications** > **Start or Configure Session**.
  2. From the Customize Communication window, select **ASCII** as your type of host and select **Link Parameters**.
  3. From the Telnet ASCII window, enter the host name or the IP address of the logical partition that is sharing its resources, and enter the port number 2301 of the logical partition that is sharing its resources. Click **OK**.
  4. If you are not using an Integrated xSeries Server, go to the next step. If you are using both AIX logical partitions and Integrated xSeries Server consoles, select **IBM i Guest Partition Consoles** from the IBM i Virtual Consoles window.
  5. From the IBM i Guest Partition Consoles window, select the logical partition to which you want to connect as the console.
  6. Enter the IBM i service tools ID and password to connect to the AIX logical partition.
- If you use Telnet, connect to the virtual console using Telnet from an MS-DOS command prompt.
  1. From an MS-DOS command prompt, use the Telnet command to connect to your server and port 2301 (`telnet` *xxxxxx* `2301`).
  2. If you are not using an Integrated xSeries Server, go to the next step. If you are using both AIX logical partitions and Integrated xSeries Server consoles, select **IBM i Guest Partition Consoles** from the IBM i Virtual Consoles window.

3. From the IBM i Guest Partition Consoles window, select the logical partition to which you want to connect as the console.
4. Enter the IBM i service tools ID and password to connect to the AIX logical partition.

Starting the network-server description for an AIX logical partition

You can start the network-server description (NWSD) for an AIX logical partition that uses IBM i resources to make the resources defined in the NWSD available to the AIX logical partition.

About this task

To start (vary on) the NWSD for an AIX logical partition, complete the following tasks:

Procedure

1. If you use IBM Navigator for i, start the NWSD by using IBM Navigator for i.
   a) Click **Network** > **Windows Administration** > **Integrated xSeries Servers**
   b) Right-click the name of the NWSD that you want to start.
   c) Click **Start**.
2. If you use the character-based interface, start the NWSD by using the character-based interface:
   a) Type `WRKCFGSTS *NWS` and press Enter.
   b) Type **1** next to the NWSD that you want to start and press Enter.

Creating an IBM i logical partition that uses IBM i virtual I/O resources

You can create an IBM i logical partition that uses IBM i virtual I/O resources on servers that are managed by a Hardware Management Console (HMC). This allows you to maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Before you begin

To set up an IBM i logical partition that uses IBM i virtual I/O resources, configure the following items:

- You must create a virtual server SCSI adapter for the IBM i logical partition that provides virtual SCSI disk resources, and create a virtual client SCSI adapter for the IBM i that uses virtual SCSI disk resources. You can then set up the IBM i logical partition with the virtual server SCSI adapter to provide disk resources to the IBM i logical partition with the virtual client SCSI adapter through the virtual SCSI connection.
- You can create a virtual serial connection between the IBM i logical partition that provides the virtual resources and the IBM i logical partition that uses the virtual resources. A virtual serial connection allows you to connect to the IBM i logical partition that uses the virtual resources from the IBM i logical partition that provides the virtual resources.
- If you want to use virtual Ethernet, create two virtual Ethernet adapters on the IBM i logical partition that uses virtual I/O resources. Both virtual Ethernet adapters must be set to connect to a virtual Ethernet adapter on the IBM i logical partition that provides the virtual I/O resources. In other words, all three virtual Ethernet adapters must be set to the same virtual LAN ID.

Both logical partitions must use IBM i 6.1.1, or later.

Alternatively, you can create a Virtual I/O Server logical partition and configure the IBM i logical partition to use the virtual SCSI and virtual Ethernet resources of the Virtual I/O Server logical partition. You might need to enter a PowerVM Editions activation code to create a Virtual I/O Server logical partition on your server.

About this task

For more information about creating a logical partition, see “Creating logical partitions” on page 79. For more information about adding IBM i hosted virtual SCSI adapters when the HMC is at version 8.7.0, or later, see Managing virtual storage for a partition.

Creating a network-server description and a network-server storage space for an IBM i logical partition that uses IBM i resources

A *network-server description (NWSD)* is an IBM i object that describes the storage resources that are used by an integrated operating environment. An NWSD can be linked to one or more network-server storage spaces. Create an NWSD to assign storage to an IBM i logical partition that uses IBM i virtual I/O resources.

Before you begin

If you assign multiple NWSDs to an IBM i logical partition that uses IBM i virtual I/O resources, ensure that only one of those NWSDs is set to provide virtual optical disk resources. Restrict the optical devices on all of the other NWSDs by adding `RSTDDEVRSC(*ALLOPT)` to the `CRTNWSD` parameters for those NWSDs.

About this task

To create an NWSD and a network-server storage space for an IBM i logical partition that uses IBM i resources, follow these steps:

Procedure

1. Determine the correct SCSI server resource name.
   - If there is only one SCSI server adapter corresponding to a given client logical partition, and that adapter has its remote logical partition and remote slot configured correctly, you can specify `*AUTO` as the RSRCNAME in your NWSD.
   - Otherwise, you must determine the actual resource name. At an IBM i command line on the IBM i logical partition that provides virtual resources, type `WRKHDWRSC *CMN`, and find a controller resource with type 290B and a converged location code that corresponds to the SCSI server adapter at the Hardware Management Console (HMC). This resource name will be used later to specify the SCSI server resource.
2. At an IBM i command line on the logical partition that provides resources, type `CRTNWSD` to create a network server description, and press F4 for prompts.
3. Specify the following information:

   The default or suggested parameter values are provided within the parentheses. These settings are relevant only to a logical partition. After the installation, if your root file system (/) is not installed on the first partition of the first disk, you must set a root parameter.

   - `NWSD (Provide a name for the NWSD)`
   - `RSRCNAME (*AUTO or the resource name of the SCSI server resource)`
   - `TYPE(*GUEST *OPSYS)`
   - `ONLINE (*NO or *YES)`
   - `PARTITION ('Provide the name of your IBM i logical partition that uses IBM i resources')`

     As an alternative to the Partition parameter, you can also specify a logical partition number by typing `PTNNBR(`*`integer`*`)` where *integer* is the number of the logical partition you are specifying.
   - `CODEPAGE (437)`
   - `TCPPORTCFG (*NONE)`

- RSTDDEVRSC (for virtual CD and tape devices) (*NONE for the NWSD that is to provide virtual optical disc resources, *ALLOPT for all other NWSDs)
- SYNCTIME (*TYPE)
- IPLSRC (*NWSSTG)
  - You can store a kernel in a disk partition of a virtual disk (a network-server storage space (NWSSTG)). By specifying the IPLSRC (*NWSSTG) parameter, you are specifying that the IBM i logical partition will start from a disk partition on that virtual disk. The disk partition on the virtual disk must be formatted as type PReP Boot (type 0x41) and marked as a device that starts.
  - To start an NWSD with a load source from a stream file, set the IPLSRC parameter to *STMF and set the IPLSTMF parameter to point to the load source. You must have read access to the file and the path leading to the file to use the vary on command.
- IPLSTMF (*NONE)
- IPLPARM (*NONE)
- PWRCTL (*YES)
  - If you specify PWRCTL (*YES), perform the following steps:
    a. Ensure that the server adapter in the IBM i logical partition that provides the virtual resources specifies the remote logical partition and remote slot in its configuration.
    b. Ensure that the client logical partition has the IBM i logical partition that provides the virtual resources as the power-controlling logical partition in the profile.
    c. Ensure before you activate the NWSD that the partition profile for the IBM i logical partition that uses IBM i virtual I/O resources has been saved by activating the logical partition from the HMC, even if the logical partition does not activate correctly because of the absence of virtual devices.
  - If you specify PWRCTL(*NO), virtual devices will be available to the logical partition. You must shut down and restart the logical partition by using the HMC.

4. Create the network-server storage space using the interface that you prefer.

Interface	Actions
IBM Navigator for i	a. Expand **My Connections** > **your server** > **Network** > **Windows Administration**. b. Right-click the **Disk Drives** and select **New Disk**. c. In the **Disk drive name** field, specify the name that you want to give to the disk drive. d. In the **Description** field, specify a meaningful description for the disk drive. e. In the **Capacity** field, specify the size of the new disk drive in megabytes. f. Click **OK**.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides the virtual I/O resources, type the command CRTNWSSTG and press F4. The Create NWS Storage Space (CRTNWSSTG) display opens. b. In the Network-server storage space field, specify the name you want to give to the storage space. c. In the Size field, specify the size in megabytes for the new storage space. d. In the Format field, specify *OPEN. e. In the Text description field, specify a meaningful description for the storage space. f. Press Enter.

5. List the network server storage spaces on the logical partition by using the Work with Network Server Storage Spaces command.

6. If you use IBM Navigator for i, link the network-server storage space by using IBM Navigator for i.

Interface	Actions
IBM Navigator for i	a. Expand **My Connections** > **your server** > **Network** > **Windows Administration**. b. Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**. c. Select the server to which you want to link the network-server storage space. d. Select the link sequence position you want to use. e. Select one of the available data access types. f. Click **OK**.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides the virtual I/O resources, type the command ADDNWSSTGL and press F4. The Add Network-Server Storage Link (ADDNWSSTGL) display opens. b. In the Network server description field, specify the name of the network server description (NWSD). c. In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the IBM i logical partition that uses IBM i virtual I/O resources (that is, available without restarting the IBM i logical partition). d. In the Drive sequence number field, specify the link sequence position you want to use. e. Press Enter.

Connecting an IBM i logical partition that uses IBM i virtual I/O resources to the virtual console

You can connect an IBM i logical partition that uses IBM i virtual I/O resources to the virtual console so that you can install the operating system or access the command line interface for the IBM i logical partition that uses IBM i virtual I/O resources.

About this task

The virtual console provides the console function for an IBM i logical partition that uses IBM i virtual I/O resources. It is used primarily during the initial installation of the operating system. The virtual console can also be used to view server errors or to restore communication to the LAN. This console connection is used prior to configuring TCP/IP.

To open a terminal window, complete the following steps :

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Actions** > **Console** > **Open Dedicated 5250 Console**.

Starting the network-server description for an IBM i logical partition that uses IBM i virtual I/O resources

You can start the network-server description (NWSD) for an IBM i logical partition that uses IBM i resources to make the resources defined in the NWSD available to the IBM i logical partition that uses IBM i resources.

About this task

To start (vary on) the NWSD for an IBM i logical partition that uses IBM i resources, complete the following tasks:

Procedure

1. If you use IBM Navigator for i, start the NWSD by using IBM Navigator for i.
   a) Click **Network** > **Windows Administration** > **Integrated xSeries Servers**
   b) Right-click the name of the NWSD that you want to start.
   c) Click **Start**.
2. If you use the character-based interface, start the NWSD by using the character-based interface:
   a) Type `WRKCFGSTS *NWS` and press Enter.
   b) Type **1** next to the NWSD that you want to start and press Enter.

Creating a Linux logical partition that uses IBM i virtual I/O resources

You can create a Linux logical partition that uses IBM i virtual I/O resources on servers that are managed by a Hardware Management Console (HMC). This allows you to maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Before you begin

To set this up, you must create virtual SCSI adapters that connect the logical partitions with each other. You can then set up the IBM i logical partition to provide disk resources to the Linux logical partition through the virtual SCSI connection. You can also create a virtual serial connection between the IBM i logical partition and the Linux logical partition. A virtual serial connection allows you to connect to the Linux logical partition from the IBM i logical partition.

Alternatively, you can create a Virtual I/O Server logical partition and configure the Linux logical partition to use the virtual SCSI and virtual Ethernet resources of the Virtual I/O Server logical partition. You might need to enter a PowerVM Editions activation code to create a Virtual I/O Server logical partition on your server.

For more information about creating a logical partition, see “Creating logical partitions” on page 79. For more information about adding IBM i hosted virtual SCSI adapters when the HMC is at version 8.7.0, or later, see Managing virtual storage for a partition.

Creating an NWSD and a network-server storage space for a Linux logical partition

A *network-server description (NWSD)* is an IBM i object that describes the storage resources that are used by an integrated operating environment. An NWSD can be linked to one or more network-server storage spaces. Create an NWSD to assign storage to a Linux logical partition that uses IBM i resources.

About this task

To create an NWSD and a network-server storage space for a Linux logical partition that uses IBM i resources, follow these steps:

Procedure

1. Determine the correct SCSI server resource name.

    - If there is only one SCSI server adapter corresponding to a given client logical partition, and that adapter has its remote logical partition and remote slot configured correctly, you can specify *AUTO as the RSRCNAME in your NWSD.
    - Otherwise, you must determine the actual resource name. At an IBM i command line, type `WRKHDWRSC *CMN`, and find a controller resource with type 290B and a converged location code that corresponds to the SCSI server adapter at the Hardware Management Console (HMC). This resource name will be used later to specify the SCSI server resource.
2. At an IBM i command line on the logical partition that shares resources, type `CRTNWSD` and press F4 for prompts.
3. Specify the following information.

   The default or suggested parameter values are provided within the parentheses. These settings are relevant only to a logical partition. After the installation, if your root file system (/) is not installed on the first partition of the first disk, you must set a root parameter.

   - `NWSD (Provide a name for the NWSD)`
   - `RSRCNAME (*AUTO or the resource name of the SCSI server resource`
   - `TYPE(*GUEST)`
   - `ONLINE (*NO or *YES)`
   - `PARTITION ('Provide the name of your Linux logical partition')`

     As an alternative to the Partition parameter, you can also specify a logical partition number by typing `PTNNBR(`*`integer`*`)` where *integer* is the number of the logical partition you are specifying.
   - `CODEPAGE (437)`
   - `TCPPORTCFG (*NONE)`
   - `RSTDDEVRSC (for virtual CD and tape devices) (*NONE)`
   - `SYNCTIME (*TYPE)`
   - `IPLSRC (*NWSSTG)`
     - You can store a kernel in a disk partition of a virtual disk (a network-server storage space (NWSSTG)). By specifying the `IPLSRC (*NWSSTG)` parameter, you are specifying that the Linux logical partition will start from a disk partition on that virtual disk. The disk partition on the virtual disk must be formatted as type `PReP Boot` (type 0x41) and marked as a device that starts. You can format a disk partition as type `PReP Boot` by using the Linux **fdisk** command with the -t option. You can specify that the disk partition starts by using the **fdisk** command with the -a option.
     - To start an NWSD with a kernel from a stream file, set the `IPLSRC` parameter to `*STMF` and set the `IPLSTMF` parameter to point to the kernel. You must have read access to the file and the path leading to the file to use the vary on command. This value only loads the kernel. After the kernel is running, it must find a root file system. In an initial installation, the root file system might be a RAM disk that is physically attached to the kernel.
   - `IPLSTMF (*NONE)`
   - `IPLPARM (*NONE)`
   - `PWRCTL (*YES)`
     - If you specify `PWRCTL (*YES)`, perform the following steps:
       a. Ensure that the server adapter in the IBM i logical partition specifies the remote logical partition and remote slot in its configuration.
       b. Ensure that the client logical partition has the IBM i logical partition as the power-controlling logical partition in the profile.
       c. Ensure before you activate the NWSD that the client logical partition's profile has been saved to the server by activating the logical partition from the HMC, even if the client operating system does not activate correctly because of the absence of virtual devices.

- If you specify `PWRCTL(*NO)`, virtual devices will be available to the logical partition. You must shut down and restart the logical partition using the HMC.

4. If you use IBM Navigator for i, create the network-server storage space by using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration**.
   b) Right-click the **Disk Drives** and select **New Disk**.
   c) In the **Disk drive name** field, specify the name that you want to give to the disk drive.
   d) In the **Description** field, specify a meaningful description for the disk drive.
   e) In the **Capacity** field, specify the size of the new disk drive in megabytes.

      Refer to your preferred Linux distributor installation documentation to determine the size you want to use.
   f) Click **OK**.
   g) Continue with step "6" on page 103.
5. If you use a character-based interface, create the network-server storage space using the character-based interface:
   a) At an IBM i command line, type the command CRTNWSSTG and press F4.

      The Create NWS Storage Space (CRTNWSSTG) display opens.
   b) In the Network-server storage space field, specify the name you want to give to the storage space.
   c) In the Size field, specify the size in megabytes for the new storage space.

      Refer to your preferred Linux distributor installation documentation to determine the size you want to use.
   d) In the Text description field, specify a meaningful description for the storage space.
   e) Press Enter.
   f) Continue with step "7" on page 103.
6. If you use IBM Navigator for i, link the network-server storage space by using IBM Navigator for i:
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration**.
   b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
   c) Select the server to which you want to link the network-server storage space.
   d) Select the link sequence position you want to use.
   e) Select one of the available data access types.
   f) Click **OK**.

   The procedure is complete. Do not complete step "7" on page 103.
7. If you use a character-based interface, link the network-server storage space by using the character-based interface:
   a) At an IBM i command line, type the command ADDNWSSTGL and press F4.

      The Add Network-Server Storage Link (ADDNWSSTGL) display opens.
   b) In the Network server description field, specify the name of the network server description (NWSD).
   c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the Linux logical partition).
   d) In the Drive sequence number field, specify the link sequence position you want to use. If you want the system to find the next available position for you, specify *CALC.
   e) Press Enter.

Connecting to the virtual console for a Linux logical partition

You can connect to the virtual console for a Linux logical partition so that you can install the operating system or access the command line interface for the Linux logical partition.

Before you begin

You must have one of the following privileges to use the Linux virtual console.

- Remote Panel
- System Partitions - Administration

About this task

The virtual console provides the console function for a Linux server. It is used primarily during the initial installation of the operating system. The virtual console can also be used to view server errors or to restore communication to the LAN. This console connection is used prior to configuring TCP/IP.

Any Telnet client can be used as the Linux console. Multiple Telnet clients can share access to the same virtual console. To connect to a console, use Telnet to connect to port 2301 of the logical partition that is sharing its resources. TCP/IP must be configured and running on at least one IBM i logical partition. Complete one of the following procedures:

- If you use IBM Personal Communications, connect to a virtual console by using IBM Personal Communications.
  1. Click **Start** > **IBM Personal Communications** > **Start or Configure Session**.
  2. From the Customize Communication window, select **ASCII** as your type of host and select **Link Parameters**.
  3. From the Telnet ASCII window, enter the host name or the IP address of the logical partition that is sharing its resources, and enter port number 2301 of the logical partition sharing its resources. Click **OK**.
  4. If you are not using an Integrated xSeries Server, go to the next step. If you are using both Linux logical partitions and Integrated xSeries Server consoles, select **IBM i Guest Partition Consoles** from the IBM i Virtual Consoles window.
  5. From the IBM i Guest Partition Consoles window, select the logical partition to which you want to connect as the console.
  6. Enter the IBM i service tools ID and password to connect to the Linux logical partition.
- If you use Telnet, connect to the virtual console using Telnet from an MS-DOS command prompt.
  1. From an MS-DOS command prompt, use the Telnet command to connect to your server and port 2301 (`telnet` *xxxxxx* `2301`).
  2. If you are not using an Integrated xSeries Server, go to the next step. If you are using both Linux logical partitions and Integrated xSeries Server consoles, select **IBM i Guest Partition Consoles** from the IBM i Virtual Consoles window.
  3. From the IBM i Guest Partition Consoles window, select the logical partition to which you want to connect as the console.
  4. Enter the IBM i service tools ID and password to connect to the Linux logical partition.

Starting the network-server description for a Linux logical partition

You can start the network-server description (NWSD) for a Linux logical partition that uses IBM i resources to make the resources defined in the NWSD available to the Linux logical partition.

About this task

To start (vary on) the NWSD for a Linux logical partition, complete the following steps:

Procedure

1. If you use IBM Navigator for i, start the NWSD by using IBM Navigator for i.
   a) Click **Network** > **Windows Administration** > **Integrated xSeries Servers**
   b) Right-click the name of the NWSD that you want to start.
   c) Click **Start**.
2. If you use the character-based interface, start the NWSD by using a character-based interface:
   a) Type `WRKCFGSTS *NWS` and press Enter.
   b) Type **1** next to the NWSD that you want to start and press Enter.

Designating the service logical partition for your managed system

The *service logical partition* is the IBM i logical partition on a server that you can configure to apply server firmware updates to the service processor or to the hypervisor. You can also use the service logical partition to communicate server common hardware errors to IBM. These abilities are useful if the Hardware Management Console (HMC) is undergoing maintenance or is otherwise unable to perform these functions.

Before you begin

The preferred method for applying server firmware updates and communicating server common hardware errors to IBM is by using the HMC.

Servers that do not have IBM i logical partitions also do not have a service logical partition. If these servers are managed by an HMC, then you must use the HMC to update the server firmware, and the servers can contact service and support only through the HMC. Use a backup HMC to ensure that the servers have redundant methods for contacting service and support and for applying fixes.

You can designate only one logical partition at a time as the service logical partition for your managed system. The service logical partition for your server must be an IBM i logical partition.

Before you can designate a logical partition as the service logical partition for your managed system, you must shut down the logical partition. You must also shut down the logical partition before you remove the service logical partition designation from the logical partition. If you want to change the service logical partition from one logical partition to another logical partition, you must shut down both logical partitions before using this procedure.

Note: You must designate a service logical partition on a server only after you use the HMC to create, change, delete, copy, or activate any logical partitions on the managed system. You can set up the operating system on an unpartitioned server to contact service and support, and you can use the operating system on an unpartitioned server to apply server firmware updates.

About this task

To designate one of your logical partitions as the service logical partition, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. In the **General Settings** area, select the logical partition that you want to designate as the service logical partition.
5. Click **OK**.

Resetting the managed system to a nonpartitioned configuration

You can use the Hardware Management Console (HMC) and the Advanced System Management Interface (ASMI) to erase all of your logical partitions and reset the managed system to a nonpartitioned configuration. When you reset the managed system, all of the physical hardware resources are assigned to a single logical partition. This allows you to use the managed system as if it were a single, nonpartitioned server.

Before you begin

⚠ **Attention:** By resetting a partitioned managed system to a nonpartitioned configuration, you will lose all of your logical partition configuration data. However, resetting the managed system does not erase the operating systems and data from disk units on that managed system.

Before you reset the managed system, ensure that the hardware placement in the managed system supports a nonpartitioned configuration. If the hardware placement in the managed system does not support a nonpartitioned configuration, you must move the hardware so that the hardware placement supports a nonpartitioned configuration. For more information about how to place the hardware in your managed system to support a nonpartitioned configuration, contact your marketing representative or business partner.

Also, if you plan to use an operating system that is already installed on one of the logical partitions on the managed system (instead of reinstalling the operating system after you reset the managed system), consider how the console used by that operating system will change when you reset the managed system. If the operating system that you want to use is AIX, log into AIX and enable the login prompt for the virtual serial port vty0 using either the System Management Interface Tool (SMIT) or the **chdev** command. You can then reset the managed system, use a physical serial console to log into AIX, and use SMIT or the **chcons** command to change the console device to the console device you want to use.

You must have an ASMI login profile with an administrator authority level.

Parts of this procedure must be performed *at your HMC* (not connected remotely). Ensure that you have physical access to the HMC before you begin.

About this task

To reset a managed system with logical partitions to a nonpartitioned configuration using the HMC, follow these steps:

Procedure

1. Shut down all logical partitions on your managed system using operating system procedures.

   For more information about shutting down logical partitions using operating system procedures, see the following information:

   - For logical partitions running AIX, see "Shutting down AIX logical partitions" on page 130.
   - For logical partitions running Linux, see "Shutting down Linux logical partitions" on page 136.
   - For logical partitions running Virtual I/O Server, see "Shutting down Virtual I/O Server logical partitions by using the HMC" on page 137.
2. If the managed system powered off automatically when you shut down the last logical partition, power on the managed system to the Standby state. Complete the following steps:
   a) In the navigation pane of your HMC, open **Systems Management** and click **Servers**.
   b) In the work pane, select the managed system, click the **Tasks** button, and click **Operations** > **Power On**.
   c) Select the power-on mode of **Partition Standby** and click **OK**.
   d) Wait until the work pane displays a Standby state for the managed system.
3. Initialize the profile data on the HMC. Complete the following:

a) In the work pane, select the managed system, click the **Tasks** button, and click **Configuration** > **Manage Partition Data** > **Initialize**.

b) Click **Yes** to confirm.

4. Clear the logical partition configuration data on the managed system. Complete the following *at your HMC* (not connected remotely):

   a) In the navigation pane, click **HMC Management**.

   b) In the work pane, click **Open Restricted Shell Terminal**.

   c) Type the command: `lpcfgop -m managed_system_name -o clear`, where *managed_system_name* is the name of the managed system as it is displayed in the work pane.

   d) Enter **1** to confirm.

   This step takes several seconds to complete.

5. Optional: If you no longer intend to manage the system using the HMC, remove the connection between the HMC and the managed system. To remove the connection between the HMC and the managed system, complete the following:

   a) In the work pane, select the managed system, click the **Tasks** button, and click **Connections** > **Reset or Remove Connection**.

   b) Select **Remove connection** and click **OK**.

6. Access the Advanced System Management Interface (ASMI) using a Web browser on a PC. If you do not already have a PC that is set up to access the ASMI on the managed system, you need to set up the PC at this point.

   For instructions, see Accessing the ASMI using a web browser.

7. On the ASMI Welcome pane, log in using the admin user ID (enter `admin` into **User ID**, enter the `admin` password into **Password**, and click **Log In**).
8. In the navigation pane, expand **Power/Restart Control** and click **Power On/Off System**.
9. Set **Boot to server firmware** to `Running`.
10. Click **Save settings and power off**.
11. Click **Power On/Off System** periodically to refresh the window. Repeat this step until **Current system power state: Off** is displayed in the navigation pane.
12. Click **Save settings and power on**.
13. Wait for the managed system to restart.

    It can take several minutes for the managed system and operating system to restart completely.

Deleting a logical partition

You can use the Hardware Management Console (HMC) to delete a logical partition and all of the partition profiles associated with the logical partition.

Before you begin

You cannot delete a logical partition if it is the service logical partition of your managed system. Before you can delete such a logical partition, you must designate another logical partition as the service logical partition of your managed system or remove the service logical partition designation from the logical partition.

Before you delete a logical partition, complete the following steps:

1. Shut down the logical partition that you plan to delete. For instructions, see "Shutting down and restarting logical partitions" on page 130.
2. If the logical partition that you plan to delete is a Virtual I/O Server logical partition that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*), remove the paging VIOS partition from the shared memory pool. For instructions, see "Removing a paging VIOS partition from the shared memory pool" on page 122.

When you delete a logical partition that uses shared memory, the HMC automatically performs the following tasks:

- The HMC removes the shared memory partition from the shared memory pool.
- The HMC returns the physical memory that was allocated to the shared memory partition for its I/O devices to the shared memory pool so that the hypervisor can allocate the physical memory to other shared memory partitions.
- The HMC releases the paging space device that was allocated to the shared memory partition so that it becomes available for other shared memory partitions to use.

Attention: This procedure erases the logical partition and the logical partition configuration data stored on the partition profiles.

About this task

To delete a logical partition, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Tasks** > **Delete Partition**. You can select the **Cleanup associated Virtual I/O Server mappings** check box and the **Delete associated virtual disks** check box.
4. Click **OK** to confirm.

Configuring virtual resources for logical partitions

You can use the Hardware Management Console (HMC) to configure virtual resources such as virtual Ethernet adapters, Host Ethernet Adapter, and shared processor pools. Configure virtual resources to help optimize the use of physical system resources.

Note: HEA is not supported on POWER9 processor-based server.

Configuring Active Memory Expansion for AIX logical partitions

You can configure Active Memory Expansion for an AIX logical partition by using the Hardware Management Console (HMC). Configuring Active Memory Expansion for a logical partition compresses the memory of the logical partition and thus expands its memory capacity.

Before you begin

You can configure Active Memory Expansion for logical partitions that use dedicated memory and logical partitions that use shared memory.

Before you start, complete the following tasks:

1. Complete the required preparation tasks for Active Memory Expansion and ensure that your configuration meets the configuration requirements for Active Memory Expansion. For instructions, see “Preparing to configure Active Memory Expansion” on page 64.
2. Enter the required activation code to enable Active Memory Expansion on the server. For instructions, see “Entering the activation code for Active Memory Expansion” on page 117.

About this task

For more information about configuring Active Memory Expansion on a logical partition, see Changing memory settings.

What to do next

After you configure Active Memory Expansion for the logical partition, monitor the performance of the workload and adjust the configuration, if necessary. For instructions, see "Adjusting the Active Memory Expansion configuration to improve performance" on page 203.

Configuring a virtual Ethernet adapter

You can configure a virtual Ethernet adapter dynamically for a running logical partition by using the Hardware Management Console (HMC). Doing so will connect the logical partition to a virtual LAN (VLAN).

Before you begin

You can dynamically configure a virtual Ethernet adapter for a Linux logical partition only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

If you plan to configure an Ethernet adapter for a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you might need to adjust the amount of I/O entitled memory assigned to the shared memory partition before you configure the adapter:

- If the I/O entitled memory mode of the shared memory partition is set to the auto mode, you do not need to take action. When you configure the new Ethernet adapter, the HMC automatically increases the I/O entitled memory of the shared memory partition to accommodate the new adapter.
- If the I/O entitled memory mode of the shared memory partition is set to the manual mode, you must increase the I/O entitled memory that is assigned to the shared memory partition to accommodate the new adapter. For instructions, see "Adding and removing I/O entitled memory dynamically to and from a shared memory partition" on page 151.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see Managing virtual network connections.

Note: In the AIX operating system, the receive buffers pools of the Virtual Ethernet increases in size and shrinks. When the load increases, the receive buffers pools increases by several buffers. The buffer pools can increase in size until the maximum value is reached (defined by the *buf_mode* attribute). When the load decreases, the receive buffers pools shrink to the minimum value (defined by the *buf_mode* attribute). The *buf_mode* attribute has the following possible values:

- *min* - Allocate the minimum buffer values. Increase as required and shrink back to min values.
- *max* - Allocate the maximum buffer values. Shrink is disabled. Fail the device open operation if the maximum values cannot be allocated.
- *max_min* - Attempt the maximum mode. If the maximum values cannot be allocated, then fall back to minimum mode.

You can run the **`entstat`** command on the Virtual Ethernet adapter to display the number of buffers that are allocated by the device.

What to do next

After you have finished, access any existing partition profiles for the logical partition and add the virtual Ethernet adapters to those partition profiles. The virtual Ethernet adapter is lost if you shut down the logical partition and activate that logical partition using a partition profile that does not have the virtual Ethernet adapter in it.

Changing the VLAN IDs of a virtual Ethernet adapter

You can dynamically change the VLAN IDs of a virtual Ethernet adapter for a running logical partition by using the Hardware Management Console (HMC).

Before you begin

Before you start, verify that the Virtual I/O Server is at version 2.2.0.0, or later.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see Managing virtual network connections.

What to do next

After changing the VLAN IDs of a virtual Ethernet adapter, access any existing partition profiles for the logical partition and add the VLAN IDs of the virtual Ethernet adapters to the partition profiles. The dynamically set values in the additional VLANs are lost if you have shut down and then activated the logical partition by using a partition profile that does not have the new list of VLAN IDs of the virtual Ethernet adapter.

Configuring the Quality of Service priority for a virtual Ethernet adapter

You can dynamically configure the Quality of Service (QoS) priority of a virtual Ethernet adapter of a running logical partition by using the Hardware Management Console (HMC). You can prioritize the logical partition network traffic by specifying the value of IEEE 802.1Q priority level for each virtual Ethernet adapter.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see Managing virtual network connections.

The virtual Ethernet QoS priority level values are in the range of 1 - 7. The following table lists the different priority levels.

VLAN user priority level	Quality of Service priority
1	Background
2	Spare
0 (default)	Best effort
3	Excellent effort
4	Controlled load
5	Video < 100 ms latency and jitter
6	Voice < 10 ms latency and jitter
7	Network control

MAC address controls using the HMC

The HMC version 7 release 7.2.0 or later, introduces HMC controls and policies for MAC address assignment to virtual Ethernet adapters and to logical Host Ethernet Adapter (LHEA).

By using the HMC, you can perform the following tasks:

- Specify a custom MAC address for the virtual Ethernet adapters of a logical partition.

  Note: For a virtual Ethernet adapter, the default value is the HMC-generated MAC address.

Tip: Avoid specifying a MAC address to enable automatic generation of a MAC address.

- Apply the following controls to MAC address overrides specified at the operating system level:
  - Allow all operating-system-defined MAC addresses
  - Deny all operating-system-defined MAC addresses
  - Specify allowable operating-system-defined MAC addresses (you can specify a maximum of four operating-system defined MAC addresses)

    Note: By default, all overrides are allowed. This is applicable to both virtual Ethernet adapters and LHEA. HEA is not supported on POWER9 processor-based server.
- Specify an optional initial MAC address for a virtual Ethernet adapter to replace an HMC-generated initial MAC address.

Note: The MAC address controls can be applied only when creating a logical partition, modifying a partition profile, or dynamically adding a virtual Ethernet adapter and logical Host Ethernet Adapter. You cannot dynamically modify an existing virtual Ethernet adapter or LHEA to add or change MAC controls.

The rules for custom virtual Ethernet MAC addresses are:

- The MAC address must be 6 bytes long.
- The bit 1 of byte 0 is reserved for Ethernet multicasting and must always be off.
- The bit 2 of byte 0 indicates that the MAC address is a locally administered address and must always be on.

Configuring the MAC address controls for a virtual Ethernet adapter

By using the Hardware Management Console (HMC), you can configure the MAC address controls of a virtual Ethernet adapter of a logical partition during logical partition creation, during partition profile modification, or when dynamically adding a virtual Ethernet adapter. You can also specify controls to MAC address overrides that are specified at the operating-system level.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see Managing virtual network connections.

Configuring a virtual Fibre Channel adapter

You can configure a virtual Fibre Channel adapter dynamically for a running logical partition by using the Hardware Management Console (HMC).

Before you begin

A Linux logical partition supports the dynamic addition of virtual Fibre Channel adapters only if the DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

If you plan to configure a virtual Fibre Channel adapter for a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you might need to adjust the amount of I/O entitled memory assigned to the shared memory partition before you configure the adapter:

- If the I/O entitled memory mode of the shared memory partition is set to the auto mode, you do not need to take action. When you configure the new virtual Fibre Channel adapter, the HMC automatically increases the I/O entitled memory of the shared memory partition to accommodate the new adapter.
- If the I/O entitled memory mode of the shared memory partition is set to the manual mode, you must increase the I/O entitled memory that is assigned to the shared memory partition to accommodate the new adapter. For instructions, see "Adding and removing I/O entitled memory dynamically to and from a shared memory partition" on page 151.

When you dynamically configure a virtual Fibre Channel adapter on a client logical partition that uses Virtual I/O Server resources, the virtual Fibre Channel adapter is lost when you restart the logical partition because the partition profile does not include the virtual Fibre Channel adapter. You cannot add the virtual Fibre Channel adapter to a partition profile after you dynamically configure it on the logical partition because the virtual Fibre Channel adapter that you add to the partition profile is assigned a different pair of worldwide port names (WWPNs) than the virtual Fibre Channel adapter that you dynamically configured on the logical partition. If you want to include the virtual Fibre Channel adapter in a partition profile, then do not dynamically configure the virtual Fibre Channel adapter on the logical partition. Instead, create the virtual Fibre Channel adapter in a partition profile and then start the logical partition using that partition profile. For instructions, see "Changing partition profile properties" on page 142.

About this task

For more information about assigning virtual Fibre Channel storage to a logical partition, see Assigning virtual Fibre Channel storage to a partition.

What to do next

If you created a virtual Fibre Channel adapter on a Virtual I/O Server logical partition, complete the following tasks:

1. Access any existing partition profiles for the Virtual I/O Server logical partition and add the virtual Fibre Channel adapter to those partition profiles. The virtual Fibre Channel adapter is lost when you shut down the Virtual I/O Server logical partition and activate it by using a partition profile that does not include the virtual Fibre Channel adapter.
2. Assign the virtual Fibre Channel adapter to a physical port on the physical Fibre Channel adapter that is connected to the physical storage that you want the associated client logical partition to access. For instructions, see Assigning the virtual Fibre Channel adapter to a physical Fibre Channel adapter.

Configuring physical ports on a Host Ethernet Adapter

You can use a Hardware Management Console (HMC) to configure the properties of each physical port on a Host Ethernet Adapter (HEA). These properties include port speed, duplex mode, maximum packet size, flow control setting, and the promiscuous logical partition for unicast packets. The physical port properties are also used by the logical ports that are associated with each physical port. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

About this task

For more information about managing Host Ethernet Adapters on a logical partition when the HMC is at version 8.7.0, or later, see Managing Host Ethernet Adapters.

What to do next

After this procedure is complete, you might need to reconfigure any logical ports that are associated with the changed physical ports. For example, if you change the maximum packet size on the physical port, you might also need to access the operating systems that use the resources on that physical port and change the maximum packet size for the corresponding logical ports.

Configuring shared processor pools

If your managed system supports more than one shared processor pool, you can use the Hardware Management Console (HMC) to configure shared processor pools on your managed system in addition to the default shared processor pool. These additional shared processor pools allow you to limit the processor usage of the logical partitions that belong to the shared processor pools. All shared processor pools other than the default shared processor pool must be configured before you can assign logical partitions to these shared processor pools.

Before you begin

You can use this procedure only if the managed system supports more than one shared processor pool and the HMC is at version 7 release 3.2.0, or later.

The default shared processor pool is preconfigured, and you cannot change the properties of the default shared processor pool.

About this task

For more information about managing shared processor pools, see Managing shared processor pools.

What to do next

After this procedure is complete, assign logical partitions to the configured shared processor pools. You can assign a logical partition to a shared processor pool at the time that you create the logical partition, or you can reassign existing logical partitions from their current shared processor pools to the shared processor pools that you configured using this procedure. For instructions, see "Reassigning logical partitions to shared processor pools" on page 118.

When you no longer want to use a shared processor pool, you can unconfigure the shared processor pool by using this procedure to set the maximum number of processing units and reserved number of processing units to 0. Before you can unconfigure a shared processor pool, you must reassign all logical partitions that use the shared processor pool to other shared processor pools.

Configuring the shared memory pool

You can configure the size of the shared memory pool, assign paging space devices to the shared memory pool, and assign one or two Virtual I/O Server (VIOS) logical partitions (that provide access to the paging space devices) to the shared memory pool using the Hardware Management Console (HMC).

Before you begin

Before you start, complete the following tasks:

1. Enter the activation code for the PowerVM Enterprise Edition. For instructions, see Entering the activation code for PowerVM Editions using the HMC version 7. The ability to share memory among multiple logical partitions is known as the PowerVM Active Memory Sharing technology. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code.
2. Ensure that your configuration meets the configuration requirements for shared memory. To review the requirements, see "Configuration requirements for shared memory" on page 65.
3. Complete the required preparation tasks. For instructions, see "Preparing to configure shared memory" on page 72.
4. Create the VIOS logical partitions (hereafter referred to as *paging VIOS partitions*) that you plan to assign to the shared memory pool, and then install the VIOS. For instructions, see "Creating additional logical partitions" on page 80 and Installing the VIOS and client logical partitions.
5. Create and configure the paging space devices that are owned by the paging VIOS partitions that you plan to assign to the shared memory pool. If you plan to use logical volumes as paging space devices, then create the logical volumes. For instructions, see "Creating a virtual disk for a VIOS logical partition by using the HMC" on page 115.
6. Verify that the HMC is at version 7 release 3.4.2, or later. For instructions, see Upgrading your HMC software.
7. Ensure that you are a super administrator or operator of the HMC.

About this task

If there is not enough physical memory available in the system to allocate to the shared memory pool, you can release to the hypervisor the physical memory that is currently assigned to logical partitions that use dedicated memory and that are shutdown. The hypervisor can then assign the released physical memory to the shared memory pool.

For more information about managing shared memory pools, see Managing shared memory pools.

What to do next

After you create the shared memory pool, you can create logical partitions that use shared memory. For instructions, see “Creating additional logical partitions” on page 80.

Creating a logical Host Ethernet Adapter for a running logical partition

If your managed system has a Host Ethernet Adapter (HEA), you can set up a logical partition to use HEA resources by using the Hardware Management Console (HMC) to create a logical Host Ethernet Adapter (LHEA) for the logical partition. A *logical Host Ethernet Adapter (LHEA)* is a representation of a physical HEA on a logical partition. An LHEA allows the logical partition to connect to external networks directly through the HEA. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Before you begin

You can add an LHEA dynamically to a running Linux logical partition only if you install Red Hat Enterprise Linux version 5.1, Red Hat Enterprise Linux version 4.6, or a later version of Red Hat Enterprise Linux on the logical partition. To add an LHEA to a Linux logical partition with a distribution other than these distributions, you must shut down the logical partition and reactivate the logical partition using a partition profile that specifies the LHEA.

If a logical partition is not currently running, you can create an LHEA for the logical partition by changing the partition profiles for the logical partition.

About this task

For more information about Logical host Ethernet adapter (LHEA) settings, see Logical host Ethernet adapter (LHEA) settings.

Results

When you are done, one or more new Ethernet adapters will be visible to the operating system of the logical partition.

Creating a virtual switch

You can create a virtual switch on a server by using the Hardware Management Console (HMC).

About this task

For more information about creating a virtual switch, see Adding a virtual network by creating a virtual network bridge.

Changing the virtual switch mode setting

When the virtual switch is created, the default setting is the Virtual Ethernet Bridging (VEB) mode. You can change the virtual switch mode to Virtual Ethernet Port Aggregation (VEPA) by using the Hardware Management Console (HMC).

About this task

For more information about changing a virtual switch, see Changing a virtual switch.

Synchronizing the virtual switch mode

When a Virtual I/O Server (VIOS) logical partition is in the shutdown state during the activation of a logical partition or when the external switch is downgraded, the Virtual Station Interface (VSI) profile type information is not updated in the VIOS.

About this task

When any of the VIOS logical partitions that are servicing the virtual switch or when the adjacent connected virtual switches are not in the current switching mode, you must synchronize the switching mode. You can synchronize the switching mode by using the Hardware Management Console (HMC).

You can use the **`chhwres`** command to synchronize the virtual switch mode.

Creating a Shared Ethernet adapter for a VIOS logical partition by using the HMC

You can create a Shared Ethernet Adapter on the Virtual I/O Server (VIOS) logical partition by using the Hardware Management Console.

About this task

To create a Shared Ethernet Adapter, be sure you meet the following requirements:

- The Hardware Management Console (HMC) must be at version 7 release 3.4.2 or later.
- Ensure the VIOS has one or more physical network devices or Logical Host Ethernet Adapters assigned to the logical partition. If a Logical Host Ethernet Adapter is assigned, the VIOS partition must be configured as the promiscuous logical partition for the Host Ethernet Adapter.
- Ensure a virtual Ethernet adapter is created on the VIOS. For instructions, see Configuring a virtual Ethernet adapter using the HMC.
- If the physical Ethernet adapter that you want to use as the shared adapter has TCP/IP configured, the VIOS must be at version 2.1.1.0 or later. If TCP/IP is not configured, the VIOS can be at any version.
- Ensure that there is a resource monitoring and control connection between the HMC and the VIOS.

Note: If you are using a prior release of the HMC or a prior version of a VIOS (with TCP/IP configured for the virtual Ethernet adapter), see Configuring virtual Ethernet on the Virtual I/O Server to create a Shared Ethernet Adapter by using the VIOS command-line interface.

For more information about adding a virtual network, see Managing Virtual Networks.

Creating a virtual disk for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to create a virtual disk on your managed system. Virtual disks are also known as *logical volumes*.

About this task

To modify virtual storage, be sure you meet the following requirements:

- The HMC must be at version 7.7.4, or later.
- The Virtual I/O Server (VIOS) must be at version 2.2.1.0, or later.
- Ensure that there is a resource monitoring and control connection between the HMC and the VIOS to manage storage.

To create a virtual disk, complete the following steps in the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. In the Virtual Storage Management page, click the **Virtual Disks** tab and click **Create Virtual Disk**.
6. Enter a virtual disk name, select a storage pool or a shared storage pool, and enter the size for the new virtual disk. If you select a shared storage pool, also specify whether you want to use thick or thin storage. By default, the storage type is thin storage. You can optionally assign the disk to a logical partition.
7. Click **OK**

   The HMC creates the new virtual disk with your specifications, and the Virtual Disks page is displayed.

   Tip: If possible, do not create virtual disks within the *rootvg* storage pool. Create one or more additional storage pools and create the virtual disks using the additional storage pools.
8. Repeat this procedure for each virtual disk that you want to create.
9. To view or change the properties of virtual disks that you created, see “Changing a virtual disk for a VIOS logical partition by using the HMC” on page 163.

What to do next

These steps are equivalent to using the **mkbdsp** command in the command-line interface.

If there is not enough disk space for the virtual disk, increase the size of the storage pool. For instructions, see “Changing a storage pool for a VIOS logical partition by using the HMC” on page 165

Creating storage pools

You can use the Hardware Management Console to create a volume-group-based or file-based storage pool on your managed system.

About this task

To create a volume-group-based storage pool, you must assign at least one physical volume to the storage pool. When you assign physical volumes to a storage pool, the Virtual I/O Server erases the information on the physical volumes, divides the physical volumes into physical partitions, and adds the capacity of the physical partitions to the storage pool. Do not add a physical volume to the storage pool if the physical volume contains data that you want to preserve.

To create storage pools, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.
- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.

To create a storage pool, complete the following steps in the Hardware Management Console:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.

3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page lists the VIOS partitions in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. In the Virtual Storage Management page, click the **Storage Pools** tab and click **Create Storage Pool**.
6. Enter a name for the storage pool and select the storage pool type.
7. Enter or select the information required to create the volume-group-based or file-based storage pool, and click **OK** to return to the Storage Pools page.

   Note: The new storage pool appears in the table. If you select one or more physical volumes that might belong to a different volume group, the Hardware Management Console displays a warning message to indicate that adding them to the new storage pool can result in data loss. To create the new storage pool with the selected physical volumes, select the Force option, and click **OK**.

Entering the activation code for Active Memory Expansion

You can enable Active Memory Expansion for a server by entering an activation code on the Hardware Management Console (HMC). When you enable Active Memory Expansion, you can configure the logical partitions that run on the server to compress their memory and thus expand their memory capacities.

Before you begin

Before you start, complete the following prerequisite tasks:

1. Complete the required preparation tasks for Active Memory Expansion and ensure that your configuration meets the configuration requirements for Active Memory Expansion. For instructions, see "Preparing to configure Active Memory Expansion" on page 64.
2. Verify that you have an activation code. You can obtain an activation code from your IBM Sales Representative or from the Capacity on Demand website. To obtain an activation code from the Capacity on Demand website, complete the following steps:

   a. Go to: http://www-912.ibm.com/pod/pod.

   b. Enter the system type and serial number of the server for which you need the activation code.

   c. Record the activation code that is displayed on the website.

About this task

To enter the activation code for Active Memory Expansion, complete the following steps by using the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system on which you plan to use Active Memory Expansion and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. In the **Capacity on Demand** area, click **CoD Functions**.
5. In the Capacity On Demand Functions page, click **Enter CoD Code**.
6. Enter the activation code and click **OK**.
7. From the **Tasks** menu, click **Properties**.

   The server Properties window is displayed.
8. Click the **Capabilities** tab.
9. Verify that the **Active Memory Expansion Capable** capability is now set to **True**.

   If the capability is set to False, then Active Memory Expansion is not enabled on the server. Obtain a valid activation code to enable Active Memory Expansion on the server.

10. Click **OK**.

What to do next

After you enable Active Memory Expansion on the server, you can configure logical partitions to use Active Memory Expansion. For instructions, see “Configuring Active Memory Expansion for AIX logical partitions” on page 108.

Reassigning logical partitions to shared processor pools

If you use more than one shared processor pool on your managed system, you can use the Hardware Management Console (HMC) to reassign logical partitions from one shared processor pool to another shared processor pool on your managed system.

Before you begin

You can use this procedure only if the managed system supports more than one shared processor pool and the HMC is at version 7 release 3.2.0, or later.

Any shared processor pool other than the default shared processor pool must be configured before you can assign a logical partition to the shared processor pool. (The default shared processor pool is preconfigured.) For instructions, see “Configuring shared processor pools” on page 112.

The HMC never allows the sum of the number of reserved processing units for a shared processor pool and the total number of processing units committed to the logical partitions that use the shared processor pool to be greater than the maximum number of processing units for the shared processor pool. (The default shared processor pool has no configured maximum number of processing units. The maximum number of processors available to the default shared processor pool is the total number of active, licensed processors on the managed system minus the number of processors that are assigned to dedicated processor partitions that are set not to share their dedicated processors.)

A shared processor pool cannot contain logical partitions that belong to different workload management groups. You therefore cannot reassign a logical partition with a defined workload management group to a shared processor pool that contains logical partitions that belong to another workload management group. (However, you can reassign a logical partition with a defined workload management group to a shared processor pool that contains only logical partitions that do not have a defined workload management group or that have the same workload management group as the reassigned logical partition.)

About this task

To reassign logical partitions from one shared processor pool to another shared processor pool by using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. In the **In the PowerVM** area, click **Shared Processor Pool**.
5. In the Shared Processor Pool page, click the **Partitions** tab.
6. Click the name of a logical partition that you want to reassign from one shared processor pool to another shared processor pool.
7. Select the new shared processor pool for the logical partition in the **Pool name (ID)** field and click **OK**.
8. Repeat steps 6 and 7 for any other logical partitions that you want to reassign from one shared processor pool to another shared processor pool.

9. Click **OK**.

Managing the shared memory pool

By using the Hardware Management Console (HMC), you can change the configuration of the shared memory pool. For example, you can change the amount of physical memory assigned to the shared memory pool, change the Virtual I/O Server logical partitions that are assigned to the shared memory pool, and add or remove paging space devices to or from the shared memory pool.

Changing the size of the shared memory pool

You can increase or decrease the amount of physical memory assigned to the shared memory pool by using the Hardware Management Console (HMC).

Before you begin

You must be a super administrator of operator of the HMC to change the size of the shared memory pool.

About this task

If there is not enough physical memory available in the system *by which to increase* the amount of memory assigned to the shared memory pool, you can release to the hypervisor the physical memory that is currently assigned to dedicated memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory pool.

If the shared memory pool has insufficient physical memory *by which to decrease* the amount of memory in the shared memory pool, you can release to the hypervisor the I/O entitled memory that is currently assigned to shared memory partitions that are shut down. The hypervisor can then remove the released physical memory from the shared memory pool.

For more information about managing shared memory pools, see Managing shared memory pools.

Adding a paging VIOS partition to the shared memory pool

You can use the Hardware Management Console (HMC) to assign a second Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) to the shared memory pool.

Before you begin

Before you assign a paging VIOS partition to the shared memory pool, complete the following steps:

1. Verify that only one paging VIOS partition is currently assigned to the shared memory pool.
2. Verify that the paging VIOS partition that is currently assigned to the shared memory pool is running.
3. Verify that the VIOS logical partition that you plan to assign to the shared memory pool is running.
4. Verify that you are a super administrator or an operator of the HMC.

About this task

When you assign a paging VIOS partition to the shared memory pool and both paging VIOS partitions have access to the same paging space devices, those paging space devices become common.

For more information about managing shared memory pools, see Managing shared memory pools.

What to do next

After you assign a second paging VIOS partition to the shared memory pool, complete the following steps:

1. If no common paging space devices are assigned to the shared memory pool, assign them to the shared memory pool. For instructions, see “Adding and removing paging space devices to and from the shared memory pool” on page 125.

2. Configure the logical partitions that use shared memory to use the paging VIOS partition that you assigned to the shared memory pool. For instructions, see "Changing the paging VIOS partitions assigned to a shared memory partition" on page 167.

Changing the paging VIOS partitions assigned to the shared memory pool

You can use the Hardware Management Console (HMC) to change the Virtual I/O Server (VIOS) logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to the shared memory pool.

Before you begin

Before you change the paging VIOS partitions that are assigned to the shared memory pool, complete the following steps:

1. Shut down all of the shared memory partitions that use the paging VIOS partition that you plan to change. You must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to change) as the primary paging VIOS partition, and you must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to change) as the secondary paging VIOS partition. For instructions, see "Shutting down and restarting logical partitions" on page 130.
2. Verify that the VIOS logical partition that you plan to assign to the shared memory pool as a paging VIOS partition is running. (This is the VIOS logical partition to which you plan to change the VIOS assignment of a paging VIOS partition.)
3. Verify that you are a super administrator or an operator of the HMC.

About this task

The following table describes the situations in which you can change a paging VIOS partition.

Table 18. Changing paging VIOS partitions

State of one paging VIOS partition	State of the other paging VIOS partition	Change options
Running or shut down	None. Only one paging VIOS partition is assigned to the shared memory pool.	You can change the VIOS assignment of the paging VIOS partition. In this situation, you also need to add the paging space devices to which the changed paging VIOS partition has access.
Running	Running	You can change the VIOS assignment of one of the paging VIOS partitions. You cannot change the VIOS assignment of both paging VIOS partitions at the same time.
Running	Shut down	You can change the VIOS assignment of only the paging VIOS partition that is shut down.
Shut down	Running	You can change the VIOS assignment of only the paging VIOS partition that is shut down.

Table 18. Changing paging VIOS partitions (continued)

State of one paging VIOS partition	State of the other paging VIOS partition	Change options
Shut down	Shut down	You cannot change the VIOS assignment of either paging VIOS partition. Instead, you can remove the paging VIOS partition that you do not want to change and then change the VIOS assignment of the remaining paging VIOS partition. In this situation, you also need to add the paging space devices to which the changed paging VIOS partition has access.

When you change the VIOS assignment of a paging VIOS partition, the following configuration changes occur to the paging space devices:

- Paging space devices that were common become independent if only one paging VIOS partition can access them.
- Paging space devices that were common remain common if both paging VIOS partitions can access them. (These are paging space devices to which all three VIOS logical partitions have access. The three VIOS logical partitions are the two VIOS logical partitions that were originally assigned to the shared memory pool as paging VIOS partitions plus the VIOS logical partition that you assigned as a paging VIOS partition when you changed the VIOS assignment of a paging VIOS partition.)
- Paging space devices that were independent become common if both paging VIOS partitions can access them.

When you change the VIOS assignment of a paging VIOS partition, the HMC changes the configuration of the shared memory partitions to use the VIOS logical partition that you assigned as the paging VIOS partition. When you activate the shared memory partition, the HMC automatically reflects the name of the VIOS logical partition that you assigned as the paging VIOS partition in the partition profile. The following examples explain this automatic change in more detail:

- A shared memory partition uses only one paging VIOS partition and you change the VIOS assignment of that paging VIOS partition from VIOS_A to VIOS_B. When you activate the shared memory partition, the HMC automatically shows VIOS_B as the paging VIOS partition in the partition profile.
- Two paging VIOS partitions are assigned to the shared memory pool. VIOS_A is assigned to the shared memory pool as PVP1 and VIOS_B is assigned to the shared memory pool as PVP2. A shared memory partition uses PVP1 as the primary paging VIOS partition and PVP2 as the secondary paging VIOS partition. You change the VIOS assignment of PVP1 from VIOS_A to VIOS_C. When you activate the shared memory partition, the HMC automatically shows VIOS_C as the primary paging VIOS partition and VIOS_B as the secondary paging VIOS partition.

For more information about managing shared memory pools, see Managing shared memory pools.

What to do next

After you change the VIOS assignment of a paging VIOS partition that is assigned to the shared memory pool, complete the following steps:

1. If necessary, assign paging space devices to the shared memory pool. For instructions, see "Adding and removing paging space devices to and from the shared memory pool" on page 125. You might need to add paging space devices in the following situations:
   - You changed the VIOS assignment of the only paging VIOS partition that is assigned to the shared memory pool. The VIOS logical partition that you assigned as the paging VIOS partition has access

to different paging space devices than the VIOS logical partition that was previously assigned as the paging VIOS partition. The paging space devices to which the current paging VIOS partition has access must be assigned to the shared memory pool for the shared memory partitions to use them.

- You removed a paging VIOS partition that was shut down and then changed the VIOS assignment of the other paging VIOS partition that was also shut down. Because you removed a paging VIOS partition from the shared memory pool, you changed the VIOS assignment of the only paging VIOS partition that is assigned to the shared memory pool. The VIOS logical partition that you assigned as the paging VIOS partition has access to different paging space devices than the VIOS logical partition that was previously assigned as the paging VIOS partition. The paging space devices to which the current paging VIOS partition has access must be assigned to the shared memory pool for the shared memory partitions to use them.
- You changed the VIOS assignment of a paging VIOS partition that provided independent paging space devices to shared memory partitions. The VIOS logical partition that you assigned as the paging VIOS partition has access to different paging space devices than the VIOS logical partition that was previously assigned as the paging VIOS partition. The independent paging space devices to which the current paging VIOS partition has access must be assigned to the shared memory pool for the shared memory partitions to continue to use independent paging space devices.

2. Activate all of the shared memory partitions that you previously shut down so that your changes can take effect. For instructions, see “Activating a logical partition” on page 126.

Removing a paging VIOS partition from the shared memory pool

You can use the Hardware Management Console (HMC) to remove a Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) from the shared memory pool.

Before you begin

Before you remove a paging VIOS partition from the shared memory pool, complete the following steps:

1. Verify that two paging VIOS partitions are currently assigned to the shared memory pool.
2. Shut down all of the shared memory partitions that use the paging VIOS partition that you plan to remove. You must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to remove) as the primary paging VIOS partition, and you must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to remove) as the secondary paging VIOS partition. For instructions, see “Shutting down and restarting logical partitions” on page 130.
3. Verify that you are a super administrator or an operator of the HMC.

About this task

The following table describes the situations in which you can remove a paging VIOS partition.

Table 19. Removing paging VIOS partitions

State of one paging VIOS partition	State of the other paging VIOS partition	Removal options
Running	Running	You can remove either paging VIOS partition.
Running	Shut down	You can remove only the paging VIOS partition that is shut down.
Shut down	Running	You can remove only the paging VIOS partition that is shut down.

Table 19. Removing paging VIOS partitions (continued)

State of one paging VIOS partition	State of the other paging VIOS partition	Removal options
Shut down	Shut down	You can remove either paging VIOS partition; however, you need to reassign the paging space devices to the shared memory pool when you activate the remaining paging VIOS partition. To avoid adding the paging space devices again, you can activate one of the paging VIOS partitions and then remove the other paging VIOS partition.

When you remove a paging VIOS partition from the shared memory pool, the following configuration changes occur:

- Paging space devices that were common become independent.
- The HMC changes the configuration of each shared memory partition to use the remaining paging VIOS partition as the primary and only paging VIOS partition:
  - If a shared memory partition uses only one paging VIOS partition and you remove that paging VIOS partition, the HMC changes the configuration of the shared memory partition to use the remaining paging VIOS partition. When you activate the shared memory partition, the HMC automatically reflects the name of the current paging VIOS partition in the partition profile.

    For example, two paging VIOS partitions, VIOS_A and VIOS_B, are assigned to the shared memory pool. A shared memory partition, SMP1, uses only VIOS_A as its paging VIOS partition. You remove VIOS_A from the shared memory pool. When you activate SMP1, the HMC automatically shows VIOS_B as the primary and only paging VIOS partition in the partition profile.
  - If a shared memory partition uses two paging VIOS partitions and you remove a paging VIOS partition, the HMC changes the configuration of the shared memory partition to use the remaining paging VIOS partition as the primary and only paging VIOS partition. When you activate the shared memory partition, the HMC ignores the primary and secondary settings in the partition profile and assigns the remaining paging VIOS partition as the primary and only paging VIOS partition for the shared memory partition. If you want to save the configuration, you can update the partition profile or save the logical partition configuration to a new partition profile.

For more information about managing shared memory pools, see Managing shared memory pools.

What to do next

After you remove a paging VIOS partition from the shared memory pool, complete the following steps:

1. If you removed a paging VIOS partition that was shut down and the other paging VIOS partition was also shut down, complete the following steps:
   a. Activate the remaining paging VIOS partition. For instructions, see “Activating a logical partition” on page 126.
   b. Remove the remaining paging space devices from the shared memory pool and assign them again to the shared memory pool. Even though the paging space devices become independent when you remove a paging VIOS partition from the shared memory pool, they cannot be recognized as such until you reassign them to the shared memory pool. For instructions, see “Adding and removing paging space devices to and from the shared memory pool” on page 125.

2. If the paging VIOS partition that you removed was the only paging VIOS partition used by a shared memory partition and the remaining paging VIOS partition does not have access to an available paging space device that meets the size requirements of the shared memory partition, assign such a paging space device to the shared memory pool. For instructions, see "Adding and removing paging space devices to and from the shared memory pool" on page 125.
3. Activate all of the shared memory partitions that you previously shut down so that your changes can take effect. For instructions, see "Activating a logical partition" on page 126.

Reinstalling the Virtual I/O Server of a paging VIOS partition

When you reinstall the Virtual I/O Server (VIOS) that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*), you need to reconfigure the shared memory environment. For example, you might need to add the paging space devices again to the shared memory pool.

About this task

The paging VIOS partitions store information about the paging space devices that are assigned to a shared memory pool. The Hardware Management Console (HMC) obtains information about the paging space devices that are assigned to the shared memory pool from the paging VIOS partitions. When you reinstall the VIOS, the information about the paging space devices is lost. For the paging VIOS partitions to regain the information, you must assign the paging space devices again to the share memory pool after you reinstall the VIOS.

The following table shows the reconfiguration tasks that you must perform in the shared memory environment when you resinstall the Virtual I/O Server of a paging VIOS partition.

Table 20. Shared memory reconfiguration tasks for reinstalling the Virtual I/O Server of a paging VIOS partition

Number of paging VIOS partitions that are assigned to the shared memory pool	Number of paging VIOS partitions for which you want to reinstall the VIOS	Reconfiguration steps	Instructions
1	1	1. Shut down all logical partitions that use shared memory (hereafter referred to as *shared memory partitions*). 2. Reinstall the VIOS. 3. Add the paging space devices again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Installing the Virtual I/O Server manually 3. Adding and removing paging space devices to and from the shared memory pool

Table 20. Shared memory reconfiguration tasks for reinstalling the Virtual I/O Server of a paging VIOS partition (continued)

Number of paging VIOS partitions that are assigned to the shared memory pool	Number of paging VIOS partitions for which you want to reinstall the VIOS	Reconfiguration steps	Instructions
2	1	1. Shut down each shared memory partition that uses the paging VIOS partition (that you plan to reinstall) as the primary or secondary paging VIOS partition. 2. Remove the paging VIOS partition from the shared memory pool. 3. Reinstall the VIOS. 4. Add the paging VIOS partition again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Removing a paging VIOS partition from the shared memory pool 3. Installing the Virtual I/O Server manually 4. Adding a paging VIOS partition to the shared memory pool
2	2	1. Shut down all the shared memory partitions. 2. Reinstall the VIOS of each paging VIOS partition. 3. Add the paging space devices again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Installing the Virtual I/O Server manually 3. Adding and removing paging space devices to and from the shared memory pool

Adding and removing paging space devices to and from the shared memory pool

After you create the shared memory pool, you can add and remove paging space devices to and from the shared memory pool by using the Hardware Management Console (HMC).

Before you begin

Before you add a paging space device, complete the following tasks:

1. Configure the paging space device to the Virtual I/O Server (VIOS) logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to the shared memory pool. If you plan to use logical volumes as paging space devices, then create the logical volumes. For instructions, see “Creating a virtual disk for a VIOS logical partition by using the HMC” on page 115.
2. Verify that all paging VIOS partitions are running.

Before you remove a paging space device, complete the following tasks:

- If no logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) is using the paging space device, verify that the paging space device is inactive.

- If a shared memory partition is using the paging space device, verify that the shared memory partition is shut down.
- Verify that all paging VIOS partitions are running.

You must be a super administrator or operator of the HMC to add and remove paging space devices to and from the shared memory pool.

About this task

For more information about managing shared memory pools, see Managing shared memory pools.

Deleting the shared memory pool

If you no longer want any of the logical partitions to use shared memory, you can delete the shared memory pool by using the Hardware Management Console (HMC).

Before you begin

Before you start, remove all of the logical partitions that use shared memory (hereafter referred to as *shared memory partitions*) from the shared memory pool by completing one of the following tasks:

- Delete all of the shared memory partitions. For instructions, see "Deleting a logical partition" on page 107.
- Change all of the shared memory partitions to dedicated memory partitions. For instructions, see "Changing the memory mode of a logical partition" on page 169.

About this task

For more information about managing shared memory pools, see Managing shared memory pools.

Managing logical partitions

You can manage the configuration of your logical partitions by using the Hardware Management Console (HMC). The HMC allows you to adjust the hardware resources that are used by each logical partition.

Activating a logical partition

You must activate a logical partition before you can use the logical partition. When you use the Hardware Management Console (HMC), you can activate a logical partition based on its current configuration or you can activate a logical partition by activating a partition profile.

About this task

For more information about activating a logical partition, see Activating partitions.

Activating a partition profile

You can activate a partition profile by using the Hardware Management Console (HMC). When you activate a partition profile, you activate a logical partition. The system commits resources to the logical partition based on the configuration in the partition profile and starts the operating system or software that is installed on the logical partition.

Before you begin

When you activate a logical partition by activating a partition profile, you must select a partition profile. A *partition profile* is a record on the HMC that specifies a possible configuration for a logical partition.

If you plan to activate a logical partition that uses virtual resources provided by the Virtual I/O Server, you must first activate the Virtual I/O Server (VIOS) logical partition that provides the virtual resources.

If you plan to activate a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you must first activate at least one VIOS logical partition that meets the following criteria:

- The VIOS logical partition (hereafter referred to as a *paging VIOS partition*) must provide access to an available paging space device that meets the size requirements of the shared memory partition.
- The paging VIOS partition must be assigned to the shared memory pool.

If the shared memory partition is configured with redundant paging VIOS partitions, activate both paging VIOS partitions before you activate the shared memory partition.

When you activate a shared memory partition and the shared memory pool does not contain enough physical memory required for activation, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that you want to activate.

When the partition profile contains a cable card, the partition activation fails. You must remove the cable card from the profile before you activate the logical partition, because a slot that contains a cable card cannot be partitioned.

About this task

To activate a partition profile by using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management** > **Servers**, and click the system on which the logical partition is located.
2. In the work pane, select the logical partition that you want to activate.
3. From the Tasks menu, click **Operations** > **Activate** > **Profile**.
4. If you want to install the VIOS software as part of the activation process of a logical partition, complete the following steps:

   a) Click **Yes** as the value for the **Install Virtual I/O Server as part of activation process** field.

   b) Select the partition profile that you want to use to activate the logical partition.

   c) Click **OK**.

   The **Discovering Network Adapters** window is displayed because it might take some time to load the network adapters.

   d) On the **Install Virtual I/O Server** page, select the VIOS installation source, and complete the required fields.

   e) Click **OK**.

   The installation progress pane displays the status of the VIOS installation in the progress bar. To view the details about the progress of the installation, click the **Details** tab.

   f) Click **Close**.

   A message is displayed that the VIOS installation was successful. If you selected **NIM Server** as the installation source, the NIM installation starts after you click **Close** in the installation progress pane. To view the progress of the NIM installation from a virtual terminal, click **Popup Console**. When the NIM installation is complete, a message is displayed that the installation was successful.

   g) Click **OK**.

   Note: If the **Install Virtual I/O Server as part of activation process** option fails repeatedly and the `Installation of Virtual I/O Server failed. Please contact the system administrator` message is displayed, you must type the `installios -u` command from the HMC command line to continue with the installation.
5. If you want the HMC to open a terminal window or console session for the logical partition when the logical partition is activated, click **Open a terminal window or console session**.

Note: This option is disabled when you select **Yes** as the value for the **Install Virtual I/O Server as part of activation process** field.

6. If you want to use a keylock position, boot mode, or paging VIOS redundancy configuration that is different from the keylock position, boot mode, or paging VIOS redundancy configuration that is specified in the partition profile, complete the following steps:
   a) Click **Advanced**.
   b) Select the desired keylock position, boot mode, or paging VIOS redundancy configuration.
   c) Click **OK**.
7. Click **OK**.

   If the logical partition that you want to activate is a shared memory partition and there is not enough physical memory in the shared memory pool by which to activate the shared memory partition, the Release Memory Resources window is displayed.
8. Select shared memory partitions that are shut down until the available memory is equal to or greater than the requested memory and click **OK**.

Activating a logical partition based on its current configuration

You can use the Hardware Management Console (HMC) to activate a logical partition based on its current configuration instead of a partition profile. When you activate the logical partition, the system commits resources to the logical partition based on the current configuration of the logical partition and starts the operating system or software that is installed on the logical partition. Logical partitions start faster when activated based on their current configuration data than when activated with a partition profile.

Before you begin

You cannot activate a logical partition based on its current configuration if one of the following conditions is true:

- The state of the logical partition is such that the logical partition is not capable of starting. To activate the logical partition based on its current configuration, change the state of the logical partition such that it is capable of starting.
- There is no active partition profile associated with the logical partition. For example, a newly created logical partition that has never been activated does not have an active partition profile. This logical partition cannot be activated based on its current configuration because its current configuration has no resources. The first time you activate a logical partition, you must activate it by activating a partition profile.

If you plan to activate a logical partition that uses virtual resources provided by the Virtual I/O Server, you must first activate the Virtual I/O Server (VIOS) logical partition that provides the virtual resources.

If you plan to activate a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you must first activate at least one VIOS logical partition that meets the following criteria:

- The VIOS logical partition (hereafter referred to as a *paging VIOS partition*) must provide access to an available paging space device that meets the size requirements of the shared memory partition.
- The paging VIOS partition must be assigned to the shared memory pool.

If the shared memory partition is configured with redundant paging VIOS partitions, activate both paging VIOS partitions before you activate the shared memory partition.

When you activate a shared memory partition and the shared memory pool does not contain enough physical memory required for activation, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that you want to activate.

On a HMC that is at a version earlier to Version 7.8.0, if the resource configuration field of the partition is set to **Not configured**, activating a logical partition with current configuration results in an error. On a HMC at Version 7.8.0, or later, if the resource configuration field is set to **Not configured**, and the partition has a last valid configuration profile, then that profile is used to activate the partition.

About this task

To activate a logical partition based on its current configuration by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **View Partition Properties**.
4. Click the **Partition Actions** > **Operations** > **Activate**.
5. In the Activate page, select **Current Configuration** as the value for the **Partition Configuration** field.
6. Click **Finish**.

Viewing the resource configuration status of a logical partition

You can view the resource configuration status of a logical partition by using the Hardware Management Console (HMC).

About this task

To view the resource configuration of a logical partition by using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **View Partition Properties**.
4. Click the **General** tab

   When the **Resource Configuration** field displays **Configured**, the partition can be activated by using the current configuration profile. When the **Resource Configuration** field displays **Not Configured**, the partition is activated the by using the last valid configuration that was stored as a profile.
5. Click **OK**.

Applying a profile to a logical partition

On a Hardware Management Console (HMC) that is at Version 7 Release 7.8.0, or later, you can apply a profile to a logical partition without powering on the logical partition by using the HMC command-line interface.

Procedure

From the HMC command line, type the following command:

```
chsyscfg -r lpar -m managed system -o apply -n profile name
```

Where:

- *managed system* is the name of the server on which the logical partition is located.
- *profile name* is the name of the partition profile that applied to the logical partition.

Activating a system profile

You can activate many logical partitions at a time by using the Hardware Management Console (HMC) to activate a system profile. A *system profile* is an ordered list of partition profiles. When you activate a system profile, the managed system attempts to activate the partition profiles in the system profile in the order in which the partition profiles are listed.

About this task

Restriction: You cannot activate a system profile that contains partition profiles that specify shared memory.

To activate a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
4. In the Manage System Profiles page, select the profile from the list and click **Actions** > **Activate**.
5. Select the desired activation settings for the system profile and click **Continue**.

Shutting down and restarting logical partitions

You can shut down and restart logical partitions running on systems that are managed by a Hardware Management Console (HMC).

Shutting down and restarting AIX in a logical partition

You can shut down and restart AIX in a logical partition using the Hardware Management Console (HMC).

Shutting down AIX logical partitions

You can shut down AIX logical partitions using the Hardware Management Console (HMC).

About this task

To shut down an AIX logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
4. In the work pane, select the partition and click **Actions** > **Shutdown**.
5. Select one of the following options:

Option	Description
Operating System	The HMC issues the AIX **`shutdown`** command to shut down the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Operating System Immediate	The HMC issues the AIX **`shutdown -F`** command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option

Option	Description
	is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Delayed	The logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.
Immediate	The logical partition shuts down without any preset delay.

6. Click **OK**.

Restarting AIX logical partitions
You can restart AIX logical partitions using the Hardware Management Console (HMC). Restarting a logical partition shuts the logical partition down and then starts it again.

About this task

To restart an AIX logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Restart**.
4. Select one of the following options:

Option	Description
Operating System	The HMC issues the AIX **`shutdown -r`** command to shut down and restart the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Operating System Immediate	The HMC issues the AIX **`shutdown -F -r`** command to shut down and restart the AIX logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Immediate	The logical partition is restarted as quickly as possible, without notifying the logical partition.
Dump	The HMC initiates a main storage or system memory dump on the logical partition and restarts the logical partition after the dump.

5. Click **OK**.

Shutting down IBM i logical partitions

The correct way to shut down an IBM i logical partition safely is from an IBM i command line.

If you cannot shut down the IBM i logical partition from an IBM i command line, you can shut down the IBM i logical partition from the Shut Down Partition window on your HMC or from the remote control panel on the Operations Console. Using these methods can cause an abnormal shutdown and can result in loss of data.

Before you shut down an IBM i logical partition, you must perform all of the basic IBM i shutdown tasks. For example, all other users must be signed off of the IBM i logical partition before you can shut it down. If you shut down the IBM i logical partition without completing all of the required tasks, you can cause damage to data or cause the system to behave in unpredictable ways. For instructions, see Basic system operations.

Shutting down IBM i logical partitions by using the HMC
You can shut down IBM i logical partitions by using the Hardware Management Console (HMC).

Before you begin

Before you shut down the IBM i logical partition, complete the following tasks:

1. If an integrated server is active on the system, shut down the integrated server using IBM i options.
2. Ensure that all jobs are completed and all applications are ended.
3. Ensure that your partition profiles are updated with any dynamic partitioning resource changes that you want to keep when you restart the logical partition.

About this task

The correct way to shut down an IBM i logical partition from the HMC is to open an HMC 5250 emulator session and run the Power Down System (PWRDWNSYS) command.

To shut down an IBM i logical partition from the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Actions** > **Console** > **Open Dedicated 5250 Console**.
4. From the IBM i command line in the emulator session, type PWRDWNSYS OPTION (*CNTRLD) DELAY (600) and press Enter.

   The system will only shut down the IBM i logical partition you selected. The PWRDWNSYS command does not affect other IBM i logical partitions on your system. If you enter the PWRDWNSYS command with the RESTART(*YES) option, the operating system restarts, and the resource specifications of the logical partition remain the same. If you do not use the RESTART(*YES) option, then the logical partition shuts down completely, and other logical partitions will be able to take and use the resources that were used by the logical partition. Also, when you reactivate the logical partition using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition using dynamic partitioning are lost when you reactivate the logical partition using a partition profile. If the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system using the Partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.
5. If the PWRDWNSYS command does not work, you can use either a delayed shutdown or an immediate shutdown to shut down the IBM i logical partition.

   Attention: Using these methods can cause an abnormal shutdown and can result in loss of data.

Performing a delayed shutdown of an IBM i logical partition
You can perform a delayed shutdown of a logical partition using the Hardware Management Console (HMC). Using delayed shutdown is equivalent to using the power button on the remote control panel. Use

delayed shutdown only when you must shut down a logical partition, and the PWRDWNSYS command does not work.

Before you begin

When you use the delayed shutdown option, the logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it will end abnormally and the next restart might take a long time.

About this task

To perform a delayed shutdown of an IBM i logical partition using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Shutdown**.
4. In the Shut Down Partitions page, select **Delayed** and click **OK**.

Performing an immediate shutdown of an IBM i logical partition
When you perform the immediate shutdown option by using the Hardware Management Console (HMC), the system shuts down without any preset delay. Using immediate shutdown is equivalent to using function 8 on the remote control panel.

Before you begin

Attention: Using immediate shutdown can cause an abnormal IPL of the IBM i logical partition and possibly cause loss of data. Use immediate shutdown only when an IBM i logical partition cannot shut down using PWRDWNSYS or delayed shutdown.

About this task

To perform an immediate shutdown of an IBM i logical partition using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Shutdown**.
4. In the Shut Down Partitions page, select **Immediate** and click **OK**.

Shutting down IBM i logical partitions by using Operations Console
You can shut down IBM i logical partitions by using Operations Console.

Before you shut down the IBM i logical partition, complete the following:

1. If an integrated server is active on the system, shut down the integrated server by using IBM i options.
2. Ensure that all jobs are completed and all applications are ended.

3. Ensure that your partition profiles are updated with any dynamic partitioning resource changes that you want to keep when you restart the logical partition.

The correct way to shut down a logical partition is by using the control language (CL) command Power Down System (PWRDWNSYS).

From an IBM i command line, type PWRDWNSYS OPTION (*CNTRLD) DELAY (600) and press Enter. The system will only shut down the IBM i logical partition you selected. The PWRDWNSYS command does not affect other IBM i logical partitions on your system.

If you enter the PWRDWNSYS command with the RESTART(*YES) option, the operating system restarts, and the resource specifications of the logical partition remain the same. If you do not use the RESTART(*YES) option, then the logical partition shuts down completely, and other logical partitions will be able to take and use the resources that were used by the logical partition. Also, when you reactivate the logical partition using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition using dynamic partitioning are lost when you reactivate the logical partition using a partition profile. If the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system using the Partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

If the PWRDWNSYS command does not work, you can use the remote control panel through Operations Console to use control panel functions through a PC. The graphical user interface of the remote control panel looks similar to the physical control panel. The remote control panel installs through Operations Console. Using the remote control panel to shut down the IBM i logical partition can result in an abnormal IPL and loss of data.

Delayed shutdown

Use delayed shutdown only when you must shut down a logical partition, and when the PWRDWNSYS command does not work.

When you use the delayed shutdown option, the logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it will end abnormally and the next restart might take a long time.

Immediate shutdown

Use immediate shutdown only when an IBM i logical partition cannot shut down using PWRDWNSYS or delayed shutdown.

When you use the immediate shutdown option, the system turns off without any preset delay.

⚠ **Attention:** This might cause an abnormal IPL of the IBM i logical partition and possibly cause loss of data.

Use the remote control panel to perform a delayed shutdown or an immediate shutdown. The power button will start a delayed shutdown and function 8 will start an immediate shutdown of a system.

Restarting and shutting down IBM i in a logical partition
At times you will need to perform an initial program load (IPL) or shut down an IBM i logical partition. For example, if you want to apply a delayed fix to IBM i, you must perform an IPL before IBM i can apply the fix.

The preferred method for restarting and shutting down IBM i logical partitions is through the IBM i command line. The Hardware Management Console (HMC) does not shut down the IBM i operating system before it shuts down the logical partition. Using the HMC to restart or shut down an IBM i logical partition can result in an abnormal IPL and the loss of data. However, you might need to use the HMC to change the operating mode or IPL type of the IBM i logical partition before you restart or shut down the IBM i logical partition using the IBM i command line.

It is important to remember that, when you perform an IPL of an IBM i logical partition, you are powering off only the logical partition and not the entire managed system. Other logical partitions on your managed system continue to run when you perform an IPL on the IBM i logical partition. However, when you shut down the last logical partition that is running on a managed system, then the managed system is set to power off automatically by default. If you want, you can set the managed system properties on the HMC so that the managed system remains powered on when you shut down the last running logical partition.

Changing the operating mode for an IBM i logical partition
You can change the operating mode for an IBM i logical partition using the Hardware Management Console (HMC). The operating mode for an IBM i logical partition determines the number of options that are presented to the operator for consideration during and after the initial program load (IPL). It can also secure (lock) the control panel to prevent an unauthorized or inadvertent IPL from the control panel.

About this task

To change the IBM i operating mode of a logical partition by using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **View Partition Properties**.
4. Click the **General** tab and set **Key Lock position** to your preference, and click **OK**.

Changing the IPL type for an IBM i logical partition
When you use the Hardware Management Console (HMC) to change the IPL type, the managed system loads the Licensed Internal Code and IBM i from the location specified by the IPL type. The IPL type is also known as the IPL source, because each IPL type is associated with a different IPL source.

Before you begin

You can choose a separate IPL type for each IBM i logical partition.

⚠ **Attention:** Only use IPL type C under the direction of your service representative. Severe data loss can occur with incorrect use of this function.

About this task

To change the IBM i IPL type of a logical partition using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **View Partition Properties**.
4. Click the **General** tab and set **IPL Source** to your preference, and click **OK**.

Shutting down and restarting Linux in a logical partition

You can shut down and restart Linux logical partitions or the Linux operating system by using the Hardware Management Console (HMC).

Shutting down Linux logical partitions

You can shut down Linux logical partitions and the Linux operating system using the Hardware Management Console (HMC).

About this task

To shut down a Linux logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
4. In the work pane, select the partition and click **Actions** > **Shutdown**.
5. Select one of the following options:

Option	Description
Operating System	The HMC issues the Linux **`shutdown -h +1`** command to shut down the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Operating System Immediate	The HMC issues the Linux **`shutdown -h now`** command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Delayed	The logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.
Immediate	The logical partition shuts down without any preset delay.

6. Click **OK**.

Restarting Linux logical partitions

You can restart Linux logical partitions or the Linux operating system using the Hardware Management Console (HMC). Restarting a logical partition shuts the logical partition down and then starts it again.

About this task

To restart a Linux logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Restart**.
4. Select one of the following options:

Option	Description
Operating System	The HMC issues the Linux **`shutdown -r +1`** command to shut down and restart the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Operating System Immediate	The HMC issues the Linux **`shutdown -r now`** command to shut down and restart the logical partition as quickly as possible, bypassing messages to other users.
Immediate	The logical partition is restarted as quickly as possible, without notifying the logical partition.
Dump	The HMC allows the Linux operating system on the Linux logical partition to run a diagnostic procedure. After the diagnostic procedure is complete, the logical partition restarts. The exact diagnostic procedure depends upon which Linux operating system is installed on the logical partition and how the operating system is set. The operating system might run an OS debugger, the operating system might perform a main storage or system memory dump on the logical partition, or the operating system might not be set to run any diagnostic procedure at all.

5. Click **OK**.

Shutting down and restarting Virtual I/O Server in a logical partition

You can shut down and restart Virtual I/O Server by using the Hardware Management Console (HMC).

Shutting down Virtual I/O Server logical partitions by using the HMC

You can shut down Virtual I/O Server logical partitions by using the Hardware Management Console (HMC). You can shut down the Virtual I/O Server immediately or delay the shutdown.

Before you begin

Before you shut down the Virtual I/O Server logical partition, complete the following tasks:

- If the client logical partitions that use storage and networking virtual resources provided by the Virtual I/O Server are not configured to use virtual resources provided by a redundant Virtual I/O Server, then shut down the client logical partitions.
- Shut down each shared memory partition that accesses its paging space device using only the Virtual I/O Server logical partition that you plan to shut down. If you shut down the Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) before you shut down the shared memory partitions and a shared memory partition attempts to access memory that is located on its paging space device, the shared memory partition might fail.

  If a shared memory partition accesses its paging space device redundantly through two paging VIOS partitions, you do not need to shut down the shared memory partition. When you shut down the paging VIOS partition, the shared memory partition accesses its paging space device through the other paging VIOS partition.

About this task

To shut down a Virtual I/O Server logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.

3. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
4. In the work pane, select the partition and click **Actions** > **Shutdown**.
5. Select one of the following options:

Option	Description
Operating System	The HMC issues the Virtual I/O Server **`shutdown`** command to shut down the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Operating System Immediate	The HMC issues the Virtual I/O Server **`shutdown -force`** command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Delayed	The logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.
Immediate	The logical partition shuts down without any preset delay.

6. Click **OK**.

Restarting Virtual I/O Server logical partitions by using the HMC
You can restart Virtual I/O Server logical partitions by using the Hardware Management Console (HMC). Restarting a Virtual I/O Server logical partition shuts down the Virtual I/O Server logical partition and then starts it again.

Before you begin

Before you shut down the Virtual I/O Server logical partition, complete the following tasks:

- If the client logical partitions that use storage and networking virtual resources provided by the Virtual I/O Server are not configured to use virtual resources provided by a redundant Virtual I/O Server, then shut down the client logical partitions.
- Shut down each shared memory partition that accesses its paging space device using only the Virtual I/O Server logical partition that you plan to shut down. If you shut down the Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) before you shut down the shared memory partitions and a shared memory partition attempts to access memory that is located on its paging space device, the shared memory partition might fail.

  If a shared memory partition accesses its paging space device redundantly through two paging VIOS partitions, you do not need to shut down the shared memory partition. When you shut down the paging VIOS partition, the shared memory partition accesses its paging space device through the other paging VIOS partition.

About this task

To restart a Virtual I/O Server logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Restart**.
4. Select one of the following options:

Option	Description
Operating System	The HMC issues the Virtual I/O Server **shutdown -restart** command to shut down and restart the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Operating System Immediate	The HMC issues the Virtual I/O Server **shutdown -force -restart** command to shut down and restart the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an **Open Firmware** state.
Immediate	The logical partition is restarted as quickly as possible, without notifying the logical partition.
Dump	The HMC initiates a main storage or system memory dump on the logical partition and restarts the logical partition after the dump.

5. Click **OK**.

Results

After the Virtual I/O Server restarts, complete the following tasks:

- Activate the client logical partitions that use storage and networking virtual resources provided to them by the Virtual I/O Server.
- Activate each shared memory partition that accesses its paging space device by using only the paging VIOS partition that you restarted.

Partition time power-on

If the operating system on a logical partition is scheduled to start at a specified time, the server starts automatically if the server is not already powered on. The hosting partition on a Virtual Partition Manager managed system also starts automatically if the partitions are not already running. However, other logical partitions are not started automatically at the specified time, even if the logical partitions are set to auto start at system power-on.

Procedures for scheduling operating system power-on

Each operating system has its own procedure for scheduled power-on. For information about how to schedule power-on for a specific operating system, see the link for that operating system in the following table.

Table 21. Procedures for scheduling operating system power-on

Operating system	Procedure for scheduling operating system power-on
AIX	Run the **shutdown** command, and specify the time at which the restart must occur, by using the *-t* flag. For more information, see the shutdown command page.
IBM i	Scheduling a system shutdown and restart
Linux	Set power-on time

Hardware behavior for different system configurations

When an operating system starts at a scheduled time, the configuration of the managed system determines how the operating system starts and what is started with the operating system. The following table shows the hardware behavior for each managed system configuration.

Table 22. Hardware behavior when an operating system is scheduled to power on

System configuration	Hardware behavior when an operating system is scheduled to power on
Manufacturing default configuration (MDC), where a single partition owns all the resources of the system	The following activities occur on the server that contains the operating system that is scheduled to power on: 1. The server powers on. 2. The time power-on MDC partition starts.
Server is managed by Virtual Partition Manager	The following activities occur on the server that contains the client partition that is scheduled to power on: 1. The server powers on if the server is not already on, and the hosting partition starts if the hosting partition has not already started. 2. The time power-on partition or partitions starts. Partitions that are defined to auto start with the server power will not power on automatically when the partition that is scheduled to power on starts. To set multiple partitions to start at the same time, you must set the operating system of each partition to start then.
Server is managed by the Hardware Management Console (HMC)	The following activities occur on the server that contains the partition that is scheduled to power on: 1. The server powers on, if the server is not already on. 2. The time power-on partition or partitions starts. Partitions that are defined to auto start with the server power will not power on automatically when the partition that is scheduled to power on starts. To set multiple partitions to start at the same time, you must set the operating system of each partition to start then.

Managing partition profiles for logical partitions

You can manage the partition profiles for your logical partitions using the Hardware Management Console (HMC). You can change the resource specifications stored in your partition profiles as your needs change.

Copying a partition profile

You can create a copy of an existing partition profile using the Hardware Management Console (HMC). After you create a copy of the existing partition profile, you can change the resource allocations within the new partition profile. This allows you to create multiple, nearly identical partition profiles without having to re-enter all of the resource allocations repeatedly.

About this task

To copy a partition profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to copy and click **Actions** > **Copy**.
5. Enter the name of the new partition profile into **New profile name** and click **OK**.

Results

Viewing the vNIC properties in a partition profile

You can view the properties of a virtual Network Interface Controller (vNIC) in a partition profile by using the Hardware Management Console (HMC). You can only view the properties of the vNIC and you cannot create, edit, or remove a vNIC in a partition profile.

About this task

A virtual Network Interface Controller (vNIC) is a type of virtual Ethernet adapter that can be configured on client logical partitions. Each vNIC is backed by a single root I/O virtualization (SR-IOV) logical port that is owned by the VIOS. When the HMC is at version 8.6.0, or later, the firmware is at level FW860, or later, and the VIOS is at version 2.2.5.0, or later, a dedicated vNIC can have multiple SR-IOV logical ports on different physical ports as backing devices, and they can be hosted by the same or different Virtual I/O Servers.

To view the vNIC properties in a partition profile by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Click the partition profile name that you want to view.
5. In the Logical Partition Profile Properties page, click the **Virtual Adapters** tab. If vNICs exists in the profile, they will be displayed. Click the Adapter ID of the vNIC adapter for which you want to view the properties.

   In the **General** tab of the Virtual NIC Adapter Properties page, you can view the **Adapter ID**, **Port ID**, **Hosting VIOS**, **Failover Priority**, and **Capacity** fields. The **Backing Devices** table displays all backing devices that are available. When there is more than one backing device, the **Auto Priority Failback** field is displayed.
6. In the Virtual NIC Adapter Properties page, click the **Advanced** tab.

   You can view the **Port VLAN ID**, **PVID Priority**, **VLAN Restrictions**, **MAC Address**, and **MAC Address Restrictions** fields.

Changing partition profile properties

You can change the properties of a partition profile using the Hardware Management Console (HMC). Changing the properties of a partition profile changes the resource amounts that are assigned to a logical partition when you shut down and restart the logical partition using the changed partition profile.

Before you begin

A partition profile stores the required number of processors, memory, and hardware resources assigned to that profile. Any partition profile property changes are not applied to the logical partition until you activate the partition profile.

If you plan to change a partition profile that specifies dedicated memory to a partition profile that specifies shared memory, be aware of the following actions:

- The HMC automatically deletes all of the physical I/O adapters specified in the partition profile. You can assign only virtual adapters to logical partitions that use shared memory.
- You must specify shared processors. Logical partitions that use shared memory must also use shared processors.

About this task

To change partition profile properties using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to change and click **Actions** > **Edit**.

   To add, remove, or change the vNIC adapter settings, you can run the **`chsyscfg`** command from the HMC command line. To add vNIC backing devices to a partition or to remove vNIC backing devices from a partition, and to change the vNIC auto-failback policy or to change the vNIC backing device failover policy, run the **`chhwres`** command from the HMC command line.

   When the HMC is at Version 9.1.0, or later, you can use the *max_capacity* field in the vNIC backing device attribute of the **`chsyscfg`** command to configure vNIC backing devices. You can also use the *max_capacity* attribute of the **`chsyscfg`** command to configure a single root I/O virtualization (SR-IOV) Ethernet logical port.
5. Make the appropriate changes and click **OK**.

What to do next

If you created at least one virtual fibre channel adapter, complete the following tasks to connect the logical partition to its storage:

1. Activate the logical partition. When you activate the logical partition, the HMC assigns a pair of worldwide port names (WWPNs) to the virtual fibre channel adapter. For instructions, see “Activating a logical partition” on page 126.
2. Restart the Virtual I/O Server (that provides the connection to a physical fibre channel adapter) or run the **`syscfg`** command. This enables the Virtual I/O Server to recognize the WWPNs of the virtual fibre channel adapter on the client logical partition For instructions, see “Restarting Virtual I/O Server logical partitions by using the HMC” on page 138.
3. Assign the virtual fibre channel adapter on the client logical partition to a physical port of a physical fibre channel adapter. For instructions, see “Changing virtual Fibre Channel for a Virtual I/O Server by using the HMC” on page 166.

Related information

chsyscfg command

chhwres command

lshwres command

Deleting a partition profile

You can delete a partition profile using the HMC Hardware Management Console (HMC). This allows you to remove partition profiles that you no longer use.

Before you begin

Note: You cannot delete a partition profile that is the default partition profile for the logical partition. If the partition profile you want to delete is the default partition profile, you must first change the default profile to another partition profile.

About this task

To delete a partition profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to delete and click **Actions** > **Delete**.
5. Click **OK** to confirm.

Adding the PPT ratio to the partition profile

When the Hardware Management Console (HMC) is at Version 9.1.0, or later, you can add the Physical Page Table (PPT) ratio to the partition profile.

About this task

You must be a super administrator to complete this task and the server must be in the operating state.

You can run the **chsyscfg** command from the HMC command-line interface to add the PPT ratio to the logical partition profile. To view the PPT ratio of a logical partition profile, run the **lssyscfg** command.

Related information

lssyscfg command

chsyscfg command

Managing system profiles

You can manage the system profiles on your managed system using the Hardware Management Console (HMC). You can change the logical partitions and partition profiles specified within the system profiles as the logical partitions change on your managed system.

Copying a system profile

You can use the Hardware Management Console (HMC) to create a copy of an existing system profile. After you create a copy of the existing system profile, you can change the partition profiles that are contained within the new system profile. This allows you to create multiple, nearly identical system profiles quickly and easily.

About this task

To copy a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
4. Select the system profile and click **Actions** > **Copy**.
5. Enter the name that you want to use for the copy into **New profile name** and click **OK**.

Changing a system profile

You can change which partition profiles are included in a system profile using the Hardware Management Console (HMC).

Before you begin

Restriction: You cannot add logical partitions that use shared memory to system profiles.

About this task

To change a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
4. Select the system profile that you want to change and click **Actions** > **Edit**.
5. In the **System Profile** window, select each partition profile that you want to remove from the system profile and click **Remove**.
6. For each partition profile that you want to add to the system profile, open the logical partition to which the partition profile belongs, select the partition profile, and click **Add**.
7. Click **OK**.

Validating a system profile

When you validate a system profile, the Hardware Management Console (HMC) compares the resources defined in the system profile with the resources available on the managed system. If the system profile requires more resources than are available on the managed system, a message is displayed on the HMC.

About this task

To validate a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .

2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
4. Select the system profile and click **Validate**.
5. When validation is complete, click **OK**.

Deleting a system profile

You can delete a system profile using the Hardware Management Console (HMC). This allows you to remove system profiles that you no longer use.

Before you begin

A system profile helps you activate or change the managed system from one complete set of logical partition configurations to another.

About this task

To delete a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
4. Select the system profile and click **Actions** > **Delete**.
5. Click **Yes** to confirm.

Managing the resources of a shutdown logical partition

You can use the Hardware Management Console (HMC) command-line interface to manage the resources of a shutdown logical partition.

You can use the **`chhwres`** command to remove memory, processor, and I/O resources from a shutdown logical partition.

You can change other attributes of a shutdown logical partition by changing the logical partition profile and applying the changed profile to the logical partition. Complete the following steps from the HMC command line:

1. To change the profile of a shutdown logical partition, run the following command:

   ```
   chsyscfg -r prof -m managed system -i attributes
   ```

   Note: If the profile being changed is the last activated profile, you must use the --*force* option if synchronization of the current configuration to the profile is enabled for the logical partition.

2. To apply the changed profile to the shutdown logical partition, run the following command:

   ```
   chsyscfg -r lpar -m managed system -o apply -n profile name
   ```

Managing logical partition resources dynamically

You can use the Hardware Management Console (HMC) to add, remove, or move processor, memory, and I/O resources between running logical partitions without restarting the logical partitions or the managed system.

Dynamic Platform Optimizer

POWER7 with firmware level FW760 or later, POWER8, or POWER9 processor-based servers supports the Dynamic Platform Optimizer (DPO) function. DPO is a hypervisor function that is initiated from the Hardware Management Console (HMC). DPO rearranges logical partition processors and memory on the system to improve the affinity between processors and memory of logical partitions. When DPO is running, mobility operations that target the system that is being optimized are blocked. Also, when DPO is running, many virtualization features are blocked. When a DPO operation is in progress and you want to dynamically add, remove, or move physical memory to or from running logical partitions, you must either wait for the DPO operation to complete or manually stop the DPO operation.

To help assess when DPO might be beneficial, you can use the HMC to determine affinity scores for the system and logical partitions by using the **lsmemopt** command. An affinity score is a measure of the processor-memory affinity on the system or for a partition. The score is a number in the range 0 - 100, 0 represents the worst affinity and 100 represents perfect affinity. Based on the system configuration, a score of 100 might not be attainable. A partition that has no processor and memory resources does not have an affinity score, and none is displayed for the score on the command line, when you run the **lsmemopt** command.

In addition to manually running DPO by using the **optmem** command, you can schedule DPO operations on POWER7 with firmware level FW760 or later, POWER8, or POWER9 processor-based servers. The HMC must be at Version 7.8.0 or later. The following conditions apply to the DPO operation:

- The current server affinity score of the managed system is less than or equal to the server affinity threshold that you provided.
- The affinity delta (which is the potential score minus the current score) of the managed system is greater than or equal to the affinity delta threshold of the server that you provided.

The scheduled operation sends a DPO report after the successful completion of a DPO operation, only if it is enabled in the **HMC Notifications**.

Querying affinity scores of a logical partition

On POWER7 or POWER9 processor-based servers with firmware at level FW780, or later, the HMC provides an additional flag with the **lsmemopt** command for querying the current affinity score and the potential affinity score of a logical partition.

About this task

Procedure

1. From the HMC command line, type the following command to query the current and potential logical partition affinity scores:

   ```
   lsmemopt -m managed system -r lpar -o currscore | calcscore [-p partition-names | --id
   partition-IDs]
   [-x partition-names | --xid partition-IDs]
   ```

   where:

   - *currscore* queries the current affinity scores.
   - *calcscore* queries the current and potential affinity scores.
   - *-x partition-names* or *--xid partition-IDs* specifies the list of logical partitions or logical partition IDs that must not be affected by the optimization operation.
   - *-p partition-names* or *--id partition-IDs* specifies the list of logical partitions or logical partition IDs that must be optimized.

The following example shows a sample output of the **lsmemopt** command when the *-o currscore* parameter is specified:

```
lpar_name=x,lpar_id=1,curr_lpar_score=25
```

The following example shows a sample output of the **lsmemopt** command when the *-o calcscore* parameter is specified:

```
lpar_name=x,lpar_id=1,curr_lpar_score=25,predicted_lpar_score=100
```

2. From the HMC command line, type the following command to query the system-wide affinity scores:

```
lsmemopt -m managed system -o currscore | calcscore [-p partition-names | --id partition-
IDs]
[-x partition-names | --xid partition-IDs]
```

where:

- *currscore* queries the current affinity scores.
- *calcscore* queries the current and potential affinity scores.
- *-x partition-names* or *--xid partition-IDs* specifies the list of logical partitions or logical partition IDs that must not be affected by the optimization operation.
- *-p partition-names* or *--id partition-IDs* specifies the list of logical partitions or logical partition IDs that must be optimized.

Scheduling Dynamic Platform Optimizer operations

Scheduled operation of the Dynamic Platform Optimizer (DPO) function is supported on POWER7 or POWER9 processor-based servers with firmware at level 7.6, or later. The Hardware Management Console (HMC) must be at Version 7.8.0 or later.

About this task

To schedule DPO operations by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **Schedule Operations**.
4. On the **Options** tab, click **New**.
5. Click **Monitor/Perform Dynamic Platform Optimize**.
6. Click **OK**.
7. On the Setup a Scheduled Operation page, click the **Date and Time** tab.

   You can specify the date and time when the scheduled operation must start.
8. Click **Save**.
9. On the Setup a Scheduled Operation page, click the **Repeat** tab.

   You can specify whether the scheduled operation is a single scheduled operation or a repeated scheduled operation. You can also specify the days of the week the operation must be performed, the interval and the number of repetitions. Click **Repeat Indefinitely** to perform the operation repeatedly for an indefinite period.
10. Click **Save**.
11. Setup a Scheduled Operation page, click the **Options** tab.
    a) In the **Target of Operation** area, the system name and potential and current affinity scores are displayed.

The potential affinity score is a value in the range 0 -100 and it is queried from the HMC when the schedule operations option is selected. You can also use the **lsmemopt** command to get this value from the HMC command line. The current affinity score is a value in the range 0 -100 and it is queried from the HMC when the schedule operations option is selected. You can also use the **lsmemopt** command to get this value from the HMC command line.

b) In the **Affinity Thresholds** area, you can specify a value in the range 0 -100 for the **Server Affinity Threshold** field.

c) In the **Server Affinity Delta Threshold (Potential - Current)** field, enter a value.

d) In the **Alert/Actions** area, when the email notification is not configured on the HMC, a message is displayed that informs you to configure the email notification. Click **Configure Management Console Notifications** to configure the email notifications.

e) In the **Alert/Actions** area, when the email notification is configured on the HMC, click **Notify via an email of Server Affinity Alerts** to receive email notification alerts about DPO events.

f) In the **Perform Dynamic Platform Optimization** area, click **Automatically Perform a Dynamic Platform Optimization (DPO)** to enable automatic DPO.

⚠ **Attention:** The DPO operation might automatically run constantly if the DPO does not cause the affinity to drop below either of the user-defined threshold values. This might impact system performance and block various virtualization functions. You can avoid setting the user-defined threshold values with a small interval when the auto-DPO option is enabled.

12. Click **Save**.

Starting and stopping a Dynamic Platform Optimizer operation

You can run the **optmem** command from the Hardware Management Console (HMC) command line on POWER7 or POWER9 processor-based servers with firmware at level FW760, or later, to start a Dynamic Platform Optimizer (DPO) operation or stop a DPO operation that is currently running.

Procedure

1. From the HMC command line, type the following command to start a DPO operation:

```
 optmem -m managed-system -o start -t affinity [-p partition-names | --id partition-IDs]
[-x partition-names | --xid partition-IDs]
```

   Where:

   - *-x partition-names* or *--xid partition-IDs* specifies the list of logical partitions or logical partition IDs that must not be affected by the optimization operation.
   - *-p partition-names* or *--id partition-IDs* specifies the list of logical partitions or logical partition IDs that must be optimized

2. To stop a currently running DPO operation, complete the following steps:

   a) From the HMC command line, type the following command to list the DPO operation that is currently running:

```
lsmemopt -m managed-system
```

   b) From the HMC command line, type the following command to stop the DPO operation:

```
optmem -m managed-system -o stop [--optid ID]
```

   Where:

   - *--optid* is an optional parameter that identifies the DPO operation to be canceled.
   - *ID* is the value returned by the **lsmemopt** command.

⚠ **Attention:** Stopping a DPO operation before completion might worsen the affinity state of the system as compared to the affinity state of the system when the DPO operation was started.

Managing dedicated memory dynamically

You can add, remove, and move physical memory dynamically to and from running logical partitions that use dedicated memory by using the Hardware Management Console (HMC). This allows you to adjust the physical memory allocated to each logical partition that uses dedicated memory without having to shut down the logical partitions.

When a DPO operation is in progress and you want to dynamically add, remove, or move physical memory to, or from running logical partitions, you must either wait for the DPO operation to complete, or manually stop the DPO operation.

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

To prevent any data loss during dynamic memory movement, the system first writes any data from memory pages to disk before making the memory pages available to another logical partition. Depending on the amount of memory you have requested to move, this might take some time.

Memory in each logical partition operates within its assigned minimum and maximum values. The full amount of memory that you assign to a logical partition might not be available for the logical partition to use. Static memory overhead that is required to support the assigned maximum memory affects the reserved or hidden memory amount. This static memory overhead also influences the minimum memory size of a logical partition.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add, remove, or move physical memory are run concurrently for a logical partition, the logical partition might not have the expected amount of physical memory after the concurrent tasks are complete. The logical partition might not have the expected amount of physical memory whether you specify the amount of physical memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of physical memory to be added to, removed from, or moved to or from the logical partition.

Adding dedicated memory dynamically

You can dynamically add physical memory to a running logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to increase the physical memory available to a logical partition that uses dedicated memory without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of memory resources only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic addition of memory resources include SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

To add memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Changing the Active Memory Expansion factor for AIX logical partitions
You can dynamically change the Active Memory Expansion factor for an AIX logical partition by using the Hardware Management Console (HMC). Changing the Active Memory Expansion factor for a logical partition increases or decreases the desired degree of expanded memory capacity for the logical partition.

Before you begin

You can change the Active Memory Expansion factor for logical partitions that use dedicated memory and logical partitions that use shared memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Moving dedicated memory dynamically
You can dynamically move physical memory from one running logical partition that uses dedicated memory to another using the Hardware Management Console (HMC). This allows you to reassign physical memory directly to a logical partition that uses dedicated memory that needs additional physical memory.

Before you begin

You cannot dynamically move memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

You can dynamically move memory to a running Linux only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic movement of memory resources include Novell SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

To move memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Removing dedicated memory dynamically
You can dynamically remove physical memory from a running AIX, IBM i, or Virtual I/O Server logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to reassign the physical memory to other logical partitions that use dedicated memory.

Before you begin

You cannot dynamically remove memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Managing shared memory dynamically

You can dynamically add and remove logical memory and I/O entitled memory to and from a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC).

About this task

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add or remove shared memory are run concurrently for a logical partition, the logical partition might not have the expected amount of shared memory after the concurrent tasks are complete. The logical partition might not have the expected amount of shared memory whether you specify the amount of shared memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of shared memory to be added to or removed from the logical partition.

Adding and removing logical memory dynamically to and from a shared memory partition

You can dynamically add and remove logical memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the logical memory assigned to the shared memory partition without having to shut down the logical partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of logical memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

To dynamically add and remove logical memory to and from a running logical partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see Changing memory settings.

Adding and removing I/O entitled memory dynamically to and from a shared memory partition

You can dynamically add and remove I/O entitled memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the maximum amount of physical memory that is assigned to the shared memory partition for its I/O devices without having to shut down the shared memory partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of I/O entitled memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

You can increase the amount of I/O entitled memory that is assigned to a shared memory partition when the sum of I/O entitled memory that is assigned to all shared memory partitions in the shared memory pool is less than the size of the shared memory pool minus the required amount of reserved firmware memory. If there is not enough physical memory in the shared memory pool by which to increase the I/O entitled memory to the amount specified, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that needs more I/O entitled memory.

You can decrease the amount of I/O entitled memory that is assigned to a shared memory partition only when the shared memory partition requires less physical memory for its I/O devices than the amount of I/O entitled memory that is assigned to the shared memory partition. For example, you assign 128 MB of I/O entitled memory to a shared memory partition. The shared memory partition requires a minimum of 64 MB for its I/O devices. Thus, you can decrease the I/O entitled memory that is assigned to the shared memory partition by up to 64 MB. For instructions about how to view the assigned, minimum, optimal, and maximum I/O entitled memory used by a shared memory partition, see "Determining the I/O entitled memory for a shared memory partition" on page 212.

To dynamically add and remove I/O entitled memory to and from a running shared memory partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see Changing memory settings.

Results

If you want to later change the I/O entitled memory mode back to the auto mode so that the HMC automatically adjusts the I/O entitled memory for the shared memory partition when you add or remove virtual adapters, repeat this procedure and select **Auto**. Alternatively, you can restart the shared memory partition. When you restart a shared memory partition, the I/O entitled memory mode is set to the auto mode regardless of what the I/O entitled memory mode was set to before you restarted the shared memory partition.

Changing the Active Memory Expansion factor for AIX logical partitions

You can dynamically change the Active Memory Expansion factor for an AIX logical partition by using the Hardware Management Console (HMC). Changing the Active Memory Expansion factor for a logical partition increases or decreases the desired degree of expanded memory capacity for the logical partition.

Before you begin

You can change the Active Memory Expansion factor for logical partitions that use dedicated memory and logical partitions that use shared memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Managing processor resources dynamically

You can dynamically add, remove, and move processor resources to and from running logical partitions using the Hardware Management Console (HMC). This allows you to adjust the processor resources allocated to each logical partition without having to shut down the logical partitions.

The ability to move processor resources dynamically becomes important when you need to adjust to changing workloads. Processor resources can be moved based on the minimum and maximum values that you created for the partition profile. You can move processor resources as long as the processor resources for each logical partition remains within the range specified by the minimum and maximum values for the logical partition. If the managed system uses more than one shared processor pool, you must also ensure that the number of processors used in each shared processor pool is less than or equal to the maximum number of processing units specified for each shared processor pool.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add, remove, or move processor resources are run concurrently for a logical partition, the logical partition might not have the expected number of processor resources after the concurrent tasks are complete. The logical partition might not have the expected number of processor resources whether you specify the number of processor resources that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the number of processor resources to be added to, removed from, or moved to or from the logical partition.

Adding processor resources dynamically

You can dynamically add processor resources to a running logical partition using the Hardware Management Console (HMC). This allows you to increase the processing capacity of a running logical partition without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of processor resources only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

About this task

For more information about changing the processor settings, see Changing processor settings.

Moving processor resources dynamically

You can dynamically move processor resources from one running logical partition to another using the Hardware Management Console (HMC). This allows you to reassign processor resources directly to a logical partition that needs additional processor resources.

Before you begin

A Linux logical partition supports the dynamic movement of processor resources only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.

- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

About this task

For more information about changing the processor settings, see Changing processor settings.

Removing processor resources dynamically
You can dynamically remove processor resources from a running logical partition using the Hardware Management Console (HMC). This allows you to reassign the processor resources to other logical partitions.

Before you begin

A Linux logical partition supports the dynamic removal of processor resources only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

About this task

For more information about changing the processor settings, see Changing processor settings.

Managing physical I/O devices and slots dynamically

You can dynamically add, remove, and move physical I/O devices and slots to and from running logical partitions using the Hardware Management Console (HMC). This allows logical partitions to share infrequently used I/O devices (such as optical disk drives).

Logical partitions can have desired or required I/O devices or slots. When you specify that an I/O device or slot is desired, this means either that the I/O device or slot is meant to be shared with other logical partitions, or that the I/O device or slot is optional. When you specify that an I/O device or slot is required (or dedicated), then you cannot activate the logical partition if the I/O device or slot is unavailable or in use by another logical partition.

Note: If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.

Adding physical I/O devices and slots dynamically
You can dynamically add a physical I/O slot (and the adapter and devices that are connected to that slot) to a running logical partition using the Hardware Management Console (HMC). This allows you to add I/O capabilities to a running logical partition without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

You cannot add physical I/O devices and slots to logical partitions that use shared memory. You can assign only virtual adapters to logical partitions that use shared memory.

About this task

For more information about managing physical I/O adapters, see Managing physical I/O adapters.

Moving physical I/O devices and slots dynamically
You can dynamically move a physical I/O slot (and the adapter and devices that are connected to that slot) from one running logical partition to another using the Hardware Management Console (HMC). This allows you to share a physical I/O device, such as a DVD drive, among many logical partitions.

Before you begin

Before you begin, vary off any devices that are attached to the managed system through the physical I/O slot that you want to move. You can vary off devices by using operating system commands.

⚠ **Attention:** The dynamic movement of a physical I/O slot that controls disk drives can cause unpredictable results, such as logical partition failure or loss of data.

A Linux logical partition supports the dynamic movement of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

You cannot dynamically move physical I/O devices and slots to logical partitions that use shared memory. You can assign only virtual adapters to logical partitions that use shared memory.

About this task

For more information about managing physical I/O adapters, see Managing physical I/O adapters.

Removing physical I/O devices and slots dynamically
You can dynamically remove a physical I/O slot and the adapter and devices that are connected to that slot from a running logical partition using the Hardware Management Console (HMC). This allows you to reassign the physical I/O slot to other logical partitions.

Before you begin

Before you begin, vary off any devices that are attached to the managed system through the physical I/O slot that you want to remove. You can vary off devices using operating system commands.

⚠ **Attention:** The dynamic removal of a physical I/O slot that controls disk drives can cause unpredictable results, such as logical partition failure or loss of data.

A Linux logical partition supports the dynamic removal of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

About this task

For more information about managing physical I/O adapters, see Managing physical I/O adapters.

Managing virtual adapters dynamically

You can dynamically add and remove virtual adapters to and from running logical partitions using the Hardware Management Console (HMC).

About this task

Tasks that are related to managing the virtual adapters such as adding a virtual adapter or removing an adapter are a part of managing the logical partition. When you perform tasks on virtual storage (vSCSI, virtual Fibre Channel, virtual optical device, virtual network, virtual NICs), the virtual adapters operations are handled automatically. For more information about virtual storage, see Managing virtual storage. For more information about virtual networks, see Managing virtual networks. For more information about virtual Network Interface Controllers (vNICs), see Managing virtual Network Interface Controllers.

Note: If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.

Enabling and disabling SR-IOV logical ports

You can enable or disable single root I/O virtualization (SR-IOV) logical ports by using the Hardware Management Console (HMC). The HMC must be at Version 9.1.0, or later.

Before you begin

You must be a super administrator to complete this task.

About this task

- To disable an SR-IOV logical port, type the following command from the HMC command line:

  ```
  chhwres -m <managed-system> -r sriov --rsubtype logport -o d [-p <partition-name> | --id
  <partition-ID>]
  -a “adapter_id=<adapter-id>,logical_port_id=<logical-port-id>”
  ```

  The logical partition can be either in the running or shut down state.

- To enable an SR-IOV logical port, type the following command from the HMC command line:

  ```
  chhwres -m <managed-system> -r sriov --rsubtype logport -o e [-p <partition-name> | --id
  <partition-ID>]
   -a “adapter_id=<adapter-id>,logical_port_id=<logical-port-id>”
  ```

  The logical partition can be either in the running or shut down state.

- To verify whether SR-IOV logical ports are enabled or disabled, type the following command from the HMC command line:

  ```
  lshwres -m <managed-system> -r sriov --rsubtype logport [--filter "<filter-data>"] –level
  <type>
       [-F [<attribute-names>] [--header]]
  ```

  If the value of the *is_disabled* attribute is 0, the SR-IOV logical ports are enabled. If the value of the *is_disabled* attribute is 1, the SR-IOV logical ports are disabled.

Related information

chhwres command

lshwres command

Managing SR-IOV logical ports dynamically

You can dynamically add, edit, and remove single root I/O virtualization (SR-IOV) logical ports to and from running logical partitions by using the Hardware Management Console (HMC).

Adding a single root I/O virtualization logical port to a logical partition dynamically
You can dynamically add a single root I/O virtualization (SR-IOV) logical port to a running logical partition by using the Hardware Management Console (HMC).

About this task

For more information about SR-IOV logical port settings, see SR-IOV logical port settings.

When the HMC is at Version 9.1.0, or later, you can use the *max_capacity* attribute of the **chhwres** command to specify the maximum capacity value for the SR-IOV logical port when you are adding an SR-IOV logical port to a logical partition.

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can use the *migratable* attribute of the **chhwres** command to add a migratable SR-IOV logical port to a logical partition.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Related information

chhwres command

Viewing migratable SR-IOV logical ports and SR-IOV backup virtual devices
You can view the list of single root I/O virtualization (SR-IOV) logical ports and SR-IOV backup virtual devices that can be migrated by using the Hardware Management Console (HMC) command-line interface.

About this task

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can run the **lshwres** command from the HMC command line to view the migratable attribute of the single root I/O virtualization (SR-IOV) logical port and SR-IOV backup virtual devices of migratable SR-IOV logical ports.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

For more information about SR-IOV logical port settings, see SR-IOV logical port settings.

Related information

lshwres command

Modifying a single root I/O virtualization logical port that is assigned to a logical partition dynamically
You can modify a single root I/O virtualization (SR-IOV) logical port that is assigned to a running logical partition by using the Hardware Management Console (HMC).

About this task

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can use the **chhwres** command to change the Port VLAN ID (PVID), allowed VLANs, and allowed operating system MAC addresses of a configured and migratable SR-IOV logical port. Additionally, the HMC checks whether the backup device of the SR-IOV logical port is a virtual Network Interface Controller (vNIC). If the backup device is a vNIC, the changes that were applied to the SR-IOV logical port are also applied to the vNIC and to the backing device of the vNIC.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

For more information about SR-IOV logical port settings, see SR-IOV logical port settings.

Related information

chhwres command

Removing a single root I/O virtualization logical port from a logical partition dynamically
You can dynamically remove a single root I/O virtualization (SR-IOV) logical port from a running logical partition by using the Hardware Management Console (HMC).

About this task

For more information about SR-IOV logical port settings, see SR-IOV logical port settings.

Note:

- When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you cannot remove a backup virtual device when an SR-IOV logical port is associated with it.
- When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Creating a profile with migratable single root I/O virtualization logical ports
You can create a profile with migratable single root I/O virtualization (SR-IOV) logical ports by using the Hardware Management Console (HMC).

About this task

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can create a profile with migratable SR-IOV logical ports by using the *migratable* attribute of the **mksyscfg** command. You must also specify the *backup_veth_vnetwork* attribute for the virtual Ethernet backup devices.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

For more information about SR-IOV logical port settings, see SR-IOV logical port settings.

Related information

mksyscfg command

Recovering a migratable single root I/O virtualization logical port
When the Hardware Management Console (HMC) is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, if a virtual I/O device is specified as the backup device for a migratable single root I/O virtualization (SR-IOV) logical port, the migratable SR-IOV logical port might not be available after running the migration operation was run by using the *--migsriov 2* option of the **migrlpar** command. You can use the *recover* option of the **chhwres** command to recover the SR-IOV logical port. When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware

is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

About this task

For more information about SR-IOV logical port settings, see SR-IOV logical port settings.

Related information

migrlpar command

Managing 5250 CPW dynamically

You can dynamically add, remove, and move 5250 commercial processing workload (5250 CPW) to and from running logical partitions using the Hardware Management Console (HMC).

5250 CPW is the capacity to perform 5250 online transaction processing (5250 OLTP) tasks on IBM i logical partitions. On certain servers, you can assign a percentage of the total 5250 CPW available on the managed system to each IBM i logical partition. The ability to assign 5250 CPW to IBM i logical partitions is available only for Express Configurations and Value Editions.

5250 CPW can be moved based on the desired, minimum, and maximum percentages you created for the partition profile. The desired 5250 CPW percentage you establish is the amount of 5250 CPW that you get if you do not overcommit the available 5250 CPW. The minimum and maximum values enable you to establish a range within which you can dynamically move the 5250 CPW.

⚠ **Attention:** If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, you should change the partition profile or save the logical partition configuration to a new partition profile.

Adding 5250 CPW for IBM i logical partitions dynamically

You can dynamically add 5250 commercial processing workload (5250 CPW) to a running IBM i logical partition using the Hardware Management Console (HMC). This allows you to increase the ability of the IBM i logical partition to run 5250 online transaction processing (5250 OLTP) tasks.

Before you begin

This procedure applies only to Express Configurations and Value Editions, which provide a fixed amount of processing capability for 5250 OLTP tasks.

About this task

To add 5250 CPW to a running IBM i logical partition using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management** > **Servers**, and click the managed system on which the logical partition resides.
2. In the work pane, select the logical partition, click the **Tasks** button, and click **Dynamic Partitioning** > **Processor** > **Add or Remove**.
3. Enter the amounts of 5250 CPW that you want the logical partition to have into the **5250 CPW (percent)** field in the **Current** column.
4. Adjust the settings in the **Options** area if necessary.

   You might need to increase the value in the **Timeout (minutes)** field to allow enough time for the HMC to complete the operation. (These settings relate to how the managed system adds 5250 CPW dynamically. These settings are not retained after the addition is completed.)
5. Click **OK**.

Moving 5250 CPW for IBM i logical partitions dynamically

You can dynamically move 5250 commercial processing workload (5250 CPW) from one running IBM i logical partition to another using the Hardware Management Console (HMC). This allows you to use the limited amount of 5250 CPW that is available on your managed system efficiently.

Before you begin

This procedure applies only to Express Configurations and Value Editions, which provide a fixed amount of processing capability for 5250 online transaction processing (5250 OLTP) tasks.

About this task

To move 5250 CPW from one running IBM i logical partition to another using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management** > **Servers**, and click the managed system on which the logical partitions reside.
2. In the work pane, select the logical partition from which you want to move 5250 CPW, click the **Tasks** button, and click **Dynamic Partitioning** > **Processor** > **Move**.
3. Enter the amounts of 5250 CPW that you want to move into the **5250 CPW (percent)** field in the **To move** column.
4. Select the logical partition to which you want to move 5250 CPW in **Select Destination Partition**.
5. Adjust the settings in the **Options** area if necessary.

   You might need to increase the value in the **Timeout (minutes)** field to allow enough time for the HMC to complete the operation. (These settings relate to how the managed system moves 5250 CPW dynamically. These settings are not retained after the move is completed.)
6. Click **OK**.

Removing 5250 CPW for IBM i logical partitions dynamically

You can dynamically remove 5250 commercial processing workload (5250 CPW) dynamically from a running IBM i logical partition using the Hardware Management Console (HMC). This allows you to make 5250 CPW available for assignment to other IBM i logical partitions on the managed system.

Before you begin

This procedure applies only to Express Configurations and Value Editions, which provide a fixed amount of processing capability for 5250 online transaction processing (5250 OLTP) tasks.

About this task

To remove 5250 CPW from a running IBM i logical partition using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management** > **Servers**, and click the managed system on which the logical partition resides.
2. In the work pane, select the logical partition, click the **Tasks** button, and click **Dynamic Partitioning** > **Processor** > **Add or Remove**.
3. Enter the amounts of 5250 CPW that you want the logical partition to have into the **5250 CPW (percent)** field in the **Current** column.
4. Adjust the settings in the **Options** area if necessary.

   You might need to increase the value in the **Timeout (minutes)** field to allow enough time for the HMC to complete the operation. (These settings relate to how the managed system removes 5250 CPW dynamically. These settings are not retained after the removal is completed.)
5. Click **OK**.

Scheduling the movement of resources to and from logical partitions

You can use the Hardware Management Console (HMC) to schedule the movement of dedicated memory, logical memory, dedicated processors, shared processors, and I/O devices between running logical partitions on a managed system. This allows you to move resources between running logical partitions without user intervention.

About this task

To schedule the movement of resources to or from a running logical partition using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **View Partition Properties**.
4. Click **Partition Actions** > **Schedule Operations**.
5. Click **Options** and click **New**.
6. Select **Dynamic Reconfiguration**, and click **OK**.
7. Select the date and time on which you want the movement to occur.
8. Select the **Options** tab and select the resource type (I/O, memory, or processor), the type of movement (**Add**, **Remove**, or **Move to**), the destination logical partition (if you are moving resources to another logical partition), and the quantity (in processors or in megabytes) or the I/O slot that you want to move.

   Note: You can add or remove logical memory to or from a logical partition. You cannot move logical memory from one logical partition to another logical partition.
9. If you want the operation to be repeated, select the **Repeat** tab and specify how you want the operation to be repeated.
10. Click **Save**.
11. When the message dialog displays, click **OK** to continue.

Results

When this procedure is completed, the managed system is set to perform the dynamic partitioning task at the date and time that you specify.

Saving the logical partition configuration to a partition profile

You can save the current configuration of a logical partition to a new partition profile using the Hardware Management Console (HMC). Use this procedure if you change the configuration of a logical partition using dynamic partitioning and you do not want to lose the changes when you reactivate the logical partition. This procedure allows you to save the changed configuration to a new partition profile instead of having to enter the changed resource allocations manually.

Before you begin

You can perform this procedure at any time after you initially activate a logical partition.

About this task

You can perform this procedure on active logical partitions and on logical partitions that are shut down. In either of these cases, the HMC reads the logical configuration that is stored for the logical partition in the server firmware and saves this logical configuration to the specified partition profile. For active logical partitions, the logical configuration that is stored in the server firmware is the current logical configuration of the logical partition. For logical partitions that are shut down, the logical configuration that is stored in the server firmware is the logical configuration at the time that you shut down the logical partition. Regardless of the state of the logical partition at the time that you perform this procedure, the procedure allows you to save the dynamic partitioning changes to a partition profile and use the partition profile to reactivate the logical partition without losing those changes.

After you shut down a logical partition, other logical partitions can use the resources that were used by that logical partition when the logical partition was active. Therefore, the resources available on the managed system might not support the logical partition configuration that is stored in the server firmware for the inactive logical partition. After you save the logical configuration of a logical partition that is shut down, verify that the resources available on the managed system can support the logical partition configuration that you saved to a partition profile.

When you save the logical configuration to a new partition profile, the desired amounts of memory, processors, processing units, and virtual processors in the new partition profile are set to the current amounts from the logical configuration. The minimum and maximum amounts of memory, processors, processing units, and virtual processors in the new partition profile are set to the minimum and maximum amounts from the logical configuration. For example, you start a logical partition using a partition profile that specifies a minimum of 512 MB of dedicated memory, a maximum of 2 GB of dedicated memory, and 1 GB as the desired amount of dedicated memory. The managed system has over 1 GB of physical memory available, so the logical partition has 1 GB of physical memory when it starts. You then add 1 GB of physical memory to the logical partition for a total of 2 GB of physical memory. If you shut down the logical partition and then save the logical configuration, the resulting partition profile specifies a minimum of 512 MB of dedicated memory, a maximum of 2 GB of dedicated memory, and 2 GB as the desired amount of dedicated memory.

The physical and virtual I/O devices that are set as required in the active partition profile are saved as required devices in the new partition profile. The physical and virtual I/O devices that are set as desired in the active partition profile or that were added to the logical partition through dynamic partitioning are saved as desired devices in the new partition profile. The partition workload group on the logical partition (if any) is saved as the partition workload group on the new partition profile.

To save the current configuration of a logical partition to a new partition profile using the HMC, complete the following:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Save Current Configuration**.
4. Enter the name of the new partition profile into **New profile** and click **OK**.

What to do next

After you save the logical configuration to a new partition profile, verify that the new partition profile is set correctly. In particular, verify that the required and desired settings are set correctly for your I/O devices. By default, physical and virtual I/O devices that are added to the logical partition using dynamic partitioning are saved as desired devices in the new partition profile. If you want any of these I/O devices to be required, you must change the partition profile so that the I/O device is required.

Managing virtual resources for Virtual I/O Server logical partitions by using the HMC

Use the Hardware Management Console (HMC) to manage virtual storage that is associated with Virtual I/O Server logical partitions.

Changing a virtual disk for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to view the properties of the virtual disks on your managed system, as well as to start virtual disk management tasks.

About this task

Virtual disks are also known as logical volumes. To assign the virtual disk to a client partition, ensure that the client partition owns one or more virtual SCSI adapters and that the Virtual I/O Server (VIOS) owns corresponding virtual SCSI adapters that host the client adapter.

To change a virtual disk, be sure you meet the following requirements:

- The HMC must be at version 7.7.4, or later.
- The VIOS must be at version 2.2.1.0, or later.
- Ensure that there is a resource monitoring and control connection between the HMC and the VIOS.

To view and to change virtual disks, complete the following steps in the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. Click the **Virtual Disks** tab to display a list of virtual disks on the managed system.
6. Select the virtual disk from the table that you want to change.

   If a virtual disk is defined as a paging space device and is assigned to a shared memory pool, it is dedicated to providing this function and is no longer available for any other purpose. Consequently, such a virtual disk is not listed here.
7. From the **Select Action** menu bar of the Virtual Disks table, select the storage management task you want to perform:

   - **Properties** to view the properties of the selected virtual disks.
   - **Extend** to add storage capacity to the selected virtual disks.
   - **Delete** to delete the selected virtual disk and make the storage resources that belonged to that virtual disk available to other virtual disks.
   - **Modify assignment** to change the logical partition to which the selected virtual disk is assigned, or to set the selected virtual disk so it is not assigned to any logical partitions.

Changing an optical device for a VIOS logical partition by using the Hardware Management Console

You can use the Hardware Management Console to view and to change physical optical devices and virtual optical media.

About this task

You can add optical devices to, or remove optical devices from, any logical partition, whether or not the logical partition is active. If you remove an optical device from an active logical partition, the Hardware Management Console prompts you to confirm the removal before removing the optical device. To assign an optical device to a client partition, ensure that the client partition owns one or more virtual SCSI adapters and that the VIOS owns corresponding virtual SCSI adapters that host the client adapter.

To change virtual optical media, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.
- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.
- Verify that a virtual media library exists before you manage, create, or assign virtual optical devices.

To view and change optical devices, complete the following steps in the Hardware Management Console:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Optical Devices** tab.
7. To change the logical partition assignment for a physical optical device, complete the following steps.

   (You cannot assign a physical optical device to an IBM i logical partition. An IBM i logical partition must use virtual optical devices instead.)

   a) From the Physical Optical Devices table, select the optical device that you want to change and click **Modify assignment**.

      The Modify Physical Optical Device Assignment page is displayed.

   b) Either change the logical partition to which the optical device is assigned, or set the optical device so it is not assigned to any logical partition, and click **OK**.

      The list of optical devices reflects the changes you made.
8. To change virtual optical media, click one of the following tasks in the Virtual Optical Media section:
   - **Create/Extend Library** to extend the size of the media library.
   - **Delete Library** to delete the media library and the files within the library.
   - **Add Media** to add an optical media file to the media library and make it available for assignment to a partition.
   - **Modify partition assignment** to change the partition assignment for a media file by changing the virtual optical device to which a media file is assigned. You can assign read-only media to more than one partition.
   - **Delete** to delete the selected media files from the media library.

Changing a storage pool for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to extend a storage pool, to reduce or remove a storage pool, and to assign a storage pool as the default storage pool for the managed system.

About this task

To view and change storage pools, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.
- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.

To view and to change storage pools, complete the following steps in the Hardware Management Console:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Storage Pools** tab to display a list of storage pools defined for the managed system.
7. Select the storage pool from the table that you want to change.
8. From the **Select Action** menu bar of the Storage Pools table, select the storage management task that you want to perform:

   - **Properties** to view the properties of the selected storage pool.
   - **Extend** to add storage capacity to the selected storage pool. To extend logical volume-based storage pools, add physical volumes to the storage pool. To extend file-based storage pools, add space from the parent storage pool to the file-based storage pool.

     Note: You cannot add a physical volume to a storage pool if it is already assigned to a partition.
   - **Reduce** to reduce the size of the selected storage pool. To reduce logical volume-based storage pools, remove physical volumes from the storage pool. To reduce the file-based storage pool, the storage pool is deleted.

     ⚠ **Attention:** Reducing a storage pool that contains virtual disks could potentially destroy data stored on the virtual disks.

Changing a physical volume for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to view the properties of the physical volumes on your managed system, as well as to start physical volume management tasks.

About this task

A physical volume can be a hard disk or a logical device on a storage area network (SAN). You can either assign a physical volume directly to a logical partition, or you can add a physical volume to a storage pool and create virtual disks to assign to logical partitions from the storage pool.

To change physical volumes, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.

- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.

To view and to modify physical volumes, complete the following steps in the Hardware Management Console:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Physical Volumes** tab to display a list of physical volumes on the managed system.
7. Select the physical volume from the table that you want to change.

   If a physical volume is defined as a paging space device and is assigned to a shared memory pool, it is dedicated to providing this function and is not available for any other purpose. Consequently, such a physical volume is not listed here.
8. From the **Select Action** menu bar of the Physical Volumes table, select the storage management task that you want to perform:
   - **Properties** to view or change the properties of the selected physical volume.
   - **Modify partition assignment** to change the logical partition to which the selected physical volume is assigned, or to set the physical volume so it is not assigned to any logical partition.
   - **Add to storage pool** to add the selected physical volume to a storage pool.
   - **Remove from storage pool** to remove the selected physical volume from the selected storage pool.

Changing virtual Fibre Channel for a Virtual I/O Server by using the HMC

You can use the Hardware Management Console (HMC) to dynamically manage virtual Fibre Channel on your managed system and the partition connections for the associated physical Fibre Channel ports. Assigning one or more physical ports to a logical partition enables the partition to communicate with storage devices in a storage area network (SAN). Configuring this type of storage resource is available only when the system supports the use of virtual Fibre Channel adapters and has a physical Fibre Channel adapter installed and connected that supports N_Port ID Virtualization (NPIV) ports.

Before you begin

To assign the virtual Fibre Channel adapter to a physical port, ensure that the client logical partition owns one or more virtual Fibre Channel adapters and that the Virtual I/O Server owns corresponding virtual Fibre Channel adapters to host the client adapter.

To change a port connection assignment for a logical partition, the partition must be either in the `Not activated` or the `Running` state. If the partition is in the `Running` state, the partition must also be capable of dynamic partitioning (DLPAR).

To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client logical partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client logical partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.

To change virtual Fibre Channel, be sure you meet the following requirements:

- The HMC must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.

- Ensure that there is a resource monitoring and control connection between the HMC and the Virtual I/O Server.

About this task

To configure the physical port connections for virtual Fibre Channel, complete the following steps in the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Virtual Fibre Channel** tab.
7. Select a port with at least one available connection and click **Modify partition connections**.

   The Modify Virtual Fibre Channel Partition Assignment page is displayed.
8. Select one or more logical partitions that you want to connect to the Fibre Channel port, and click **OK**.

   Note: If you delete the client virtual Fibre Channel adapter from the partition or the partition profile, the worldwide port names associated with the port and the storage area network (SAN) are lost. If you only change the port assignment, the worldwide port names are preserved inside the partition profile. The HMC does not reuse them when it generates port names in the future. If you run out of port names, you must obtain a code key to enable an additional prefix and range of port names for use on your system.
9. Click **OK**.

   To determine the actual number of port names available on the managed system, use the HMC to view the partition properties or partition profile properties of the client logical partition.

Managing the memory configuration of a logical partition

You can use the Hardware Management Console (HMC) to change the memory configuration of a logical partition. For example, you can change the Virtual I/O Server logical partitions that are assigned to a logical partition that uses shared memory, change the memory mode of a logical partition, and dynamically add and remove dedicated or shared memory to and from a logical partition.

Changing the paging VIOS partitions assigned to a shared memory partition

You can use the Hardware Management Console (HMC) to change the primary and secondary Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to a logical partition that uses shared memory. You can also add or remove a secondary paging VIOS partition to or from a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*).

Before you begin

Before you change the paging VIOS partitions that are assigned to a shared memory partition, complete the following tasks:

1. Ensure that the Virtual I/O Server logical partitions (that you plan to assign to the shared memory partition as paging VIOS partitions) are assigned to the shared memory pool. For instructions, see "Changing the paging VIOS partitions assigned to the shared memory pool" on page 120.
2. Ensure that the paging space device (that is accessed through the paging VIOS partitions that you plan to assign to the shared memory partition) is assigned to the shared memory pool. For instructions, see "Adding and removing paging space devices to and from the shared memory pool" on page 125.

About this task

To change the paging VIOS partitions that are assigned to a shared memory partition, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to change.
5. Click **Actions**, and click **Edit**.

   The Logical Partition Profile Properties window is displayed.
6. Click the **Memory** tab.
7. Specify a Virtual I/O Server logical partition for VIOS 1 and VIOS 2.

   Table 23. Change options for the paging VIOS partitions

Desired change	Field to change
Change the Virtual I/O Server logical partition that is assigned as the primary or the only paging VIOS partition.	Select a different Virtual I/O Server logical partition for VIOS 1.
Define a secondary paging VIOS partition.	Select a Virtual I/O Server logical partition for VIOS 2.
Change the Virtual I/O Server logical partition that is assigned as the secondary paging VIOS partition.	Select a different Virtual I/O Server logical partition for VIOS 2.
Remove the secondary paging VIOS partition.	Select None for VIOS 2.

8. Click **OK**.
9. Shut down the shared memory partition and reactivate it with the changed partition profile.

What to do next

After you change the paging VIOS partitions that are assigned to a shared memory partition, restart the shared memory partition with the changed partition profile. For instructions, see “Shutting down and restarting logical partitions” on page 130.

Changing the Active Memory Expansion factor for AIX logical partitions

You can dynamically change the Active Memory Expansion factor for an AIX logical partition by using the Hardware Management Console (HMC). Changing the Active Memory Expansion factor for a logical partition increases or decreases the desired degree of expanded memory capacity for the logical partition.

Before you begin

You can change the Active Memory Expansion factor for logical partitions that use dedicated memory and logical partitions that use shared memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Changing the memory weight of a shared memory partition

You can use the Hardware Management Console (HMC) to change the memory weight of a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*). Changing the memory weight changes the probability that the shared memory partition receives physical memory from the shared memory pool in relation to other shared memory partitions.

Before you begin

A Linux shared memory partition supports changing the memory weight only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

About this task

For more information about changing memory settings on a logical partition, see Changing memory settings.

What to do next

Changing the memory weight of a shared memory partition is temporary and is not reflected in the partition profile. The new memory weight that you assigned to the shared memory partition will be lost the next time you activate the partition profile. If you want to save the changes that you made to the memory weight of the shared memory partition, either change the partition profile or save the logical partition configuration to a new partition profile.

Changing the memory mode of a logical partition

You can create multiple partition profiles for a logical partition by using the Hardware Management Console (HMC). Some of the partition profiles can specify dedicated memory and some of the partition profiles can specify shared memory. By creating partition profiles that specify both dedicated memory and shared memory for the same logical partition, you can change the memory mode of the logical partition by activating different partition profiles.

About this task

To change the memory mode of a logical partition, complete the following steps from the HMC:

Procedure

1. Create a new partition profile for the logical partition.

   For instructions, see “Creating additional partition profiles” on page 92.

   - If you plan to change a dedicated memory partition to a shared memory partition, specify the shared memory mode in the new partition profile.
   - If you plan to change a shared memory partition to a dedicated memory partition, specify the dedicated memory mode in the new partition profile.
2. Shut down the logical partition.

   For instructions, see “Shutting down and restarting logical partitions” on page 130.
3. Activate the logical partition with the new partition profile.

   For instructions, see “Activating a partition profile” on page 126.

Managing dedicated memory dynamically

You can add, remove, and move physical memory dynamically to and from running logical partitions that use dedicated memory by using the Hardware Management Console (HMC). This allows you to adjust the physical memory allocated to each logical partition that uses dedicated memory without having to shut down the logical partitions.

When a DPO operation is in progress and you want to dynamically add, remove, or move physical memory to, or from running logical partitions, you must either wait for the DPO operation to complete, or manually stop the DPO operation.

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

To prevent any data loss during dynamic memory movement, the system first writes any data from memory pages to disk before making the memory pages available to another logical partition. Depending on the amount of memory you have requested to move, this might take some time.

Memory in each logical partition operates within its assigned minimum and maximum values. The full amount of memory that you assign to a logical partition might not be available for the logical partition to use. Static memory overhead that is required to support the assigned maximum memory affects the reserved or hidden memory amount. This static memory overhead also influences the minimum memory size of a logical partition.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add, remove, or move physical memory are run concurrently for a logical partition, the logical partition might not have the expected amount of physical memory after the concurrent tasks are complete. The logical partition might not have the expected amount of physical memory whether you specify the amount of physical memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of physical memory to be added to, removed from, or moved to or from the logical partition.

Adding dedicated memory dynamically
You can dynamically add physical memory to a running logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to increase the physical memory available to a logical partition that uses dedicated memory without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of memory resources only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic addition of memory resources include SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

To add memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Moving dedicated memory dynamically
You can dynamically move physical memory from one running logical partition that uses dedicated memory to another using the Hardware Management Console (HMC). This allows you to reassign physical memory directly to a logical partition that uses dedicated memory that needs additional physical memory.

Before you begin

You cannot dynamically move memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

You can dynamically move memory to a running Linux only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic movement of memory resources include Novell SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

To move memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Removing dedicated memory dynamically
You can dynamically remove physical memory from a running AIX, IBM i, or Virtual I/O Server logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to reassign the physical memory to other logical partitions that use dedicated memory.

Before you begin

You cannot dynamically remove memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

About this task

For more information about changing the memory settings, see Changing memory settings.

Managing shared memory dynamically

You can dynamically add and remove logical memory and I/O entitled memory to and from a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC).

About this task

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.

- When dynamic logical partitioning tasks to add or remove shared memory are run concurrently for a logical partition, the logical partition might not have the expected amount of shared memory after the concurrent tasks are complete. The logical partition might not have the expected amount of shared memory whether you specify the amount of shared memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of shared memory to be added to or removed from the logical partition.

Adding and removing logical memory dynamically to and from a shared memory partition
You can dynamically add and remove logical memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the logical memory assigned to the shared memory partition without having to shut down the logical partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of logical memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

To dynamically add and remove logical memory to and from a running logical partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see Changing memory settings.

Adding and removing I/O entitled memory dynamically to and from a shared memory partition
You can dynamically add and remove I/O entitled memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the maximum amount of physical memory that is assigned to the shared memory partition for its I/O devices without having to shut down the shared memory partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of I/O entitled memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

You can increase the amount of I/O entitled memory that is assigned to a shared memory partition when the sum of I/O entitled memory that is assigned to all shared memory partitions in the shared memory pool is less than the size of the shared memory pool minus the required amount of reserved firmware memory. If there is not enough physical memory in the shared memory pool by which to increase the I/O entitled memory to the amount specified, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that needs more I/O entitled memory.

You can decrease the amount of I/O entitled memory that is assigned to a shared memory partition only when the shared memory partition requires less physical memory for its I/O devices than the amount of I/O entitled memory that is assigned to the shared memory partition. For example, you assign 128 MB of I/O entitled memory to a shared memory partition. The shared memory partition requires a minimum of 64 MB for its I/O devices. Thus, you can decrease the I/O entitled memory that is assigned to the shared memory partition by up to 64 MB. For instructions about how to view the assigned, minimum, optimal, and maximum I/O entitled memory used by a shared memory partition, see “Determining the I/O entitled memory for a shared memory partition” on page 212.

To dynamically add and remove I/O entitled memory to and from a running shared memory partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see Changing memory settings.

Results

If you want to later change the I/O entitled memory mode back to the auto mode so that the HMC automatically adjusts the I/O entitled memory for the shared memory partition when you add or remove virtual adapters, repeat this procedure and select **Auto**. Alternatively, you can restart the shared memory partition. When you restart a shared memory partition, the I/O entitled memory mode is set to the auto mode regardless of what the I/O entitled memory mode was set to before you restarted the shared memory partition.

Obtaining additional WWPNs for the server

When all of the worldwide port names (WWPNs) on the server are used, you can add more WWPNs to the server using the Hardware Management Console (HMC). Adding WWPNs allows you to create additional virtual Fibre Channel adapters on client logical partitions that use virtual resources provided by the Virtual I/O Server.

Before you begin

The server contains 32,000 pairs of WWPNs that all contain the same 6–digit prefix. Each virtual Fibre Channel adapter that you create on a client logical partition requires one pair of WWPNs. When all of the WWPNs on the server are used, you cannot create additional virtual Fibre Channel adapters on any client logical partitions until you add more WWPNs to the server. You add more WWPNs to the server by generating an activation code that contains a new WWPN prefix that contains 32,000 new pairs of WWPNs.

About this task

To obtain additional WWPNs for the server, complete the following steps from the HMC:

Procedure

1. Retrieve information about the server:
   a) In the navigation pane, click the **Resources** icon .
   b) Click **All Systems**. The **All Systems** page is displayed.
   c) In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.

      The CoD Advanced Functions Code Information window is displayed.
   d) In the **CUoD (permanent) Processor** area, click **View CUoD Code Information**.
   e) Click **Save** to save the information to a file on a remote system or to media, and click **OK**.
2. Go to the Capacity on Demand website and enter the information that you retrieved in step “1” on page 173 to generate an activation code.
3. Apply the activation code that you obtained in step “2” on page 173 to the server:
   a) In the navigation pane, click the **Resources** icon .
   b) Click **All Systems**. The **All Systems** page is displayed.

c) In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.

The CoD Advanced Functions Code Information window is displayed.

d) In the **Capacity on Demand** area, click **CoD Functions**.

e) Enter the activation code that obtained in step “2” on page 173 and click **OK**.

4. Verify that the activation code that you entered in step “3” on page 173 was applied to the server:

a) In the navigation pane, click the **Resources** icon .

b) Click **All Systems**. The **All Systems** page is displayed.

c) In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.

The CoD Advanced Functions Code Information window is displayed.

d) In the **Capacity on Demand** area, click **CoD Functions**.

e) In the Capacity On Demand Functions page, click **View CoD History Log**.

f) Verify that there is a log entry for entering the CoD advanced functions activation code and click **Close**.

What to do next

After you finish, you can create virtual Fibre Channel adapters on client logical partitions and dynamically add virtual Fibre Channel adapters to client logical partitions.

Related concepts

Virtual Fibre Channel

With N_Port ID Virtualization (NPIV), you can configure the managed system so that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter.

Setting partition-availability priorities for your managed system

To avoid shutting down mission-critical workloads when your server firmware deconfigures a failing processor, you can use the Hardware Management Console (HMC) to set partition-availablity priorities for the logical partitions on your managed system. A logical partition with a failing processor can acquire a replacement processor from logical partitions with a lower partition-availability priority. The acquisition of a replacement processor allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

About this task

To set partition-availability priorities for your managed system by using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions** > **Legacy** > **Partition Availability Priority**.
4. Select the logical partitions whose partition-availability priority you want to set, set **Availability priority** to the partition-availability priority value that you want to use for all selected logical partitions, and click **OK**.

   You can enter any value from 0 to 255 into **Availability priority**, or you can select one of the preset choices. All selected logical partitions are set to the same partition-availability priority value.

5. Repeat this procedure for other logical partitions to set the partition-availability priority for those logical partitions.

Installing new hardware for IBM i logical partitions

You can install an I/O adapter (IOA) for an IBM i logical partition.

Before you begin

When you install new hardware in an IBM i partitioned environment, you should be aware of the following things:

- Verify that your logical partition configuration is current.
- Empty positions might not be owned by a logical partition. They should be assigned to the desired logical partition before installing new adapters in them. After you install the new adapter, you must also add the adapter to the partition profile so that, when you shut down and activate the logical partition using the partition profile, the logical partition reactivates with the adapter that you added.
- A new IOA is owned by the logical partition that owns the slot, and a new device is owned by the logical partition that owns the IOA to which the device is attached.
- New processors and memory are available (unassigned) to be assigned to any logical partition.

About this task

To install an IOA for an IBM i logical partition, perform the following steps:

Procedure

1. Assign empty slots to the desired logical partition.

   For instructions, see “Managing physical I/O devices and slots dynamically” on page 154 and “Changing partition profile properties” on page 142.
2. Install the new hardware into the empty slots. For instructions, see Installing and configuring POWER9 processor-based systems and system features.

Backing up and recovering data

It is crucial that you back up your data because you never know when you might need to do a server recovery. Save everything in your system as often as possible. You might not be prepared to recover from a site loss or certain types of disk failures if you do not regularly save everything.

For more information about planning a backup and recovery strategy for the Hardware Management Console (HMC) and IBM i data, refer to the following topics:

Table 24. Backup and recovery information for the HMC and IBM i and IBM i

Topic	Description
Backing up critical HMC data	This procedure explains how to save critical HMC data (such as user information and platform-configuration files) to a backup file. This information is in the Managing the HMC topic.
Backing up partition profile data	This procedure explains how to back up the partitioning data on your HMC to a backup file on the HMC. This information is in the Managing the HMC topic.
Reinstalling the HMC machine code	This procedure explains how to reinstall the HMC interface from the recovery CD-ROM. This information is in the Managing the HMC topic.

Table 24. Backup and recovery information for the HMC and IBM i and IBM i (continued)

Topic	Description
Restoring profile data	This procedure explains how to restore the partitioning data from the backup file to the HMC. This information is in the Managing the HMC topic.
Back up your server	This information can help you develop the backup strategy for your IBM i logical partition. This information is in the Backup and recovery topic in the IBM i Knowledge Center.
Recover your server	This information can help you reload your operating system and data. This information is in the Backup and recovery topic in the IBM i Knowledge Center.

Managing logical partitions that use IBM i resources

You can manage logical partitions that use IBM i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Managing AIX logical partitions that use IBM i resources

You can manage AIX logical partitions that uses IBM i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Related information

Backup of the system image and user-defined volume groups

Installing system backups

Adding virtual disk units to an AIX logical partition

You can dynamically add virtual disk units to an AIX logical partition that uses IBM i resources. This allows you to increase the storage capacity of your AIX logical partition when needed.

About this task

Virtual disks simplify hardware configuration on the server because they do not require you to add additional physical devices to the server in order to run AIX. You can allocate up to 64 virtual disks to an AIX logical partition. Each virtual disk supports up to 1000 GB of storage. Each virtual disk appears to AIX as one actual disk unit. However, the associated space in the IBM i integrated file system is distributed across the disks that belong to the IBM i logical partition. Distributing storage across the disks provides the benefits of device parity protection through IBM i. Therefore, you do not have to use additional processing resources and memory resources by setting up device parity protection through AIX.

IBM i provides the ability to dynamically add virtual disks to an AIX logical partition. You can allocate disk space in the integrated file system and make it available to AIX without restarting the server or logical partition. The AIX administrator can also configure the newly allocated disk space and make it available without restarting the server.

To add virtual disks dynamically to an AIX logical partition, complete the following steps:

Procedure

1. If you use IBM Navigator for i, create a network-server storage space using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration** .
   b) Right-click the **Disk Drives** and select **New Disk**.
   c) In the **Disk drive name** field, specify the name that you want to give to the network-server storage space.

   d) In the **Description** field, specify a meaningful description for the network-server storage space.
   e) In the **Capacity** field, specify the size of the new network-server storage space in megabytes.

      To help you determine the size you want to use, see Installing AIX.
   f) Click **OK**.
   g) Continue with step “3” on page 177.
2. If you use a character-based interface, create a network-server storage space using the character-based interface:
   a) At an IBM i command line, type the command **CRTNWSSTG** and press F4.

      The Create NWS Storage Space (**CRTNWSSTG**) display opens.
   b) In the Network-server storage space field, specify the name you want to give to the network-server storage space.
   c) In the Size field, specify the size in megabytes for the new network-server storage space.

      To help you determine the size you want to use, see Installing AIX.
   d) In the Text description field, specify a meaningful description for the network-server storage space.
   e) Press Enter.
   f) Continue with step “4” on page 177
3. If you use IBM Navigator for i, link the network-server storage space using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration** .
   b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
   c) Select the server to which you want to link the network-server storage space.
   d) Select one of the available data access types.
   e) Click **OK**.
   f) Continue with step “5” on page 177.
4. If you use a character-based interface, link the network-server storage space using a character-based interface:
   a) At an IBM i command line, type the command **ADDNWSSTGL** and press F4.

      The Add Network-Server Storage Link (**ADDNWSSTGL**) display opens.
   b) In the Network server description field, specify the name of the network server description (NWSD).
   c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the AIX logical partition).
   d) In the Drive sequence number field, specify the link sequence position you want to use.
   e) Press Enter.
5. Activate the AIX logical partition (if it is not already activated).
6. Log in to AIX using a user name with superuser (root) privileges.
7. Configure the new virtual disk on the AIX logical partition by running the AIX command `cfgmgr`.
8. Verify that your new disk has been added and can be configured by running the AIX command `lspv`.

   When you enter `lspv` at the command prompt, the system lists the disks that are currently available to AIX.

   An example of the output for this command is below:

   ```
   # lspv
   hdisk0          00cad6aceafe8fe4                    rootvg          active
   hdisk1          none                                None
   ```

   Note the name of the new disk as it displays in the left-hand column.

9. Configure the new disk using one of the following two methods.

   - Add the new virtual disk to the root volume group by using the AIX command `extendvg rootvg` *`diskname`*, where *`diskname`* is the name of the new disk. If you use this method, you do not need to continue this procedure. You can use AIX methods to increase the file system size at a later time.
   - Create a new volume group for the new virtual disk by using the AIX command `mkvg -y` *`volgroup diskname`*, where *`volgroup`* is the name that you want to use for the new volume group and *`diskname`* is the name of the new disk.

10. Make a logical volume on the new virtual disk using the AIX `mklv -y` *`logicvol volgroup 1 diskname`* command.

    `logicvol` is the name that you want to use for the new logical volume, *`volgroup`* is the name of the new volume group, and *`diskname`* is the name of the new disk. (The numeral *`1`* indicates that the logical volume is to consist of one logical disk partition.)

11. Format the disk partition using the AIX `crfs` command.

    There are a number of optional parameters for the `crfs` command, but typically the defaults satisfy most disk uses. To format the disk partition created in the previous steps, type the following command at an AIX command prompt, where *`logicvol`* is the name of the logical volume and *`/mnt/data`* is the mount point directory at which you want to mount the new disk:

    ```
     crfs -v jfs -d logicvol -m /mnt/data
    ```

    The `crfs` command displays the following diagnostic messages:

    ```
    crfs -v jfs -d logicvol -m /mnt/data
    Based on the parameters chosen, the new /mnt/data JFS file system is limited to
    a maximum size of 134217728 (512 byte blocks)
    New File System size is 8192.
    ```

12. Verify that the mount point directory exists by using the `cd` *`/mnt/data`* command.

    `/mnt/data` is the mount point. The `crfs` command creates this directory so that you can access your new file system. If the mount point directory does not exist, then run the following command, where *`/mnt/data`* is the name of the mount point directory:

    ```
    mkdir /mnt/data
    ```

13. Verify that an entry for your new file system exists in the `/etc/filesystems` file.

    The `crfs` command automatically generates the appropriate `/etc/filesystems` entry for your new file system. To verify that the entry exists, use an AIX text editor, such as vi, to open the `/etc/filesystems` file, and look for the entry in the `/etc/filesystems` file. If the entry does not exist, use the text editor to add the entry to the `/etc/filesystems` file.

    An example of such an entry is below:

    ```
    /mnt/data:
        dev = /dev/logicvol
        vfs = jfs
        log = /dev/loglv01
        mount = true
        account = false
    ```

    This entry mounts the virtual disk every time you restart AIX.

14. Mount the virtual disk drive in the new directory by typing: `mount /dev/`*`logicvol /mnt/data`*.

    `logicvol` is the name of the logical volume and *`/mnt/data`* is the mount point directory.

Linking a network-server storage space to a network server description

You can link a network-server storage space (NWSSTG) to one or more network server descriptions (NWSDs). This allows the NWSDs and their associated logical partitions to use the data stored on the NWSSTG.

About this task

You can link an NWSSTG to an unlimited number of NWSDs. This is beneficial when multiple logical partitions need access to a single application.

When you link an NWSSTG to an NWSD, you can set up the NWSD to have read-only access to the NWSSTG, or you can set up the NWSD to read or write to the NWSSTG.

⚠ **Attention:** If more than one NWSD can write to the NWSSTG, ensure that only one NWSD can update the data at a time. Otherwise, changes made by one NWSD can be overwritten by another NWSD.

To link an NWSSTG to an NWSD, follow these steps:

Procedure

1. At an IBM i command line, type the command `ADDNWSSTGL` and press F4.
2. From the Add Server Storage Link display, provide the following information:

   ```
   NWSSTG (Name)
   NWSD (Name)
   DYNAMIC (*YES)
   DRVSEQNBR (*CALC)
   ```

3. Press F10 (Additional Parameters).
4. Enter the type of access the storage space will have.

Deleting network-server descriptions for an AIX logical partition

You can delete the IBM i network-server description (NWSD) for an AIX logical partition that uses IBM i resources. When you delete the NWSD, all the configuration information for the AIX logical partition is deleted from IBM i.

About this task

To delete the network-server description (NWSD) for an AIX logical partition, follow these steps:

Procedure

1. On an IBM i control language (CL) command line, type the command WRKNWSD and press Enter.
2. Type 8 in the Opt field to the left of the Network Server and press Enter.
3. In the Work with Configuration Status display, if the status of the NWSD is not varied off, type 2 in the Opt field to the left of the Network Server and press Enter. Otherwise, go to the next step.
4. Press F3 to return to the previous display
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

Deleting virtual disk drives for an AIX logical partition

You can delete a virtual disk drive from an AIX logical partition that uses IBM i resources to make the space available to the IBM i logical partition once more. When you delete a virtual disk drive, all of the information on the virtual disk drive is erased.

Before you begin

Before you can delete a virtual disk drive, you must unlink the virtual disk drive from the network-server description (NWSD). For instructions, see "Unlinking virtual disk drives from an AIX logical partition" on page 181.

About this task

To delete a virtual disk drive, follow these steps:

Procedure

Delete the disk drive using the interface that you prefer.

Interface	Actions
IBM Navigator for i	a. Click **Network** > **Windows Administration** > **Disk Drives**. b. Right-click the disk drive that you want to delete. c. Click **Delete** in the confirmation window.
IBM i character-based interface	a. At an IBM i control language (CL) command line, type `DLTNWSSTG` and press F4. b. Type the name of the disk drive in the Network-server storage space field and press Enter.

Using IPL types when running AIX

The IPL source (IPLSRC) parameter on the network-server description (NWSD) determines the initial program that is loaded when the NWSD is varied on. For an AIX logical partition that uses IBM i resources, the initial program is the kernel. Ensure that the IPLSRC parameter specifies the kernel location of the kernel for the AIX logical partition that uses IBM i resources.

You can set the IPLSRC parameter when you use the Create Network Server Description (CRTNWSD) command, and you can change the IPLSRC parameter when you use the Change Network Server Description (CHGNWSD) command.

Note: The IPLSRC parameter also has the values A, B, and D, which are not valid for hardware that is used by IBM i logical partitions.

The IPLSRC parameter has the following valid values.

IPLSRC values	Description
*Panel	The logical partition is started from the source indicated on the control panel.
*NWSSTG (network-server storage space)	This IPL type is used to start a logical partition from a virtual disk. The open firmware will find the kernel in the virtual disk. The open firmware searches the first virtual disk connected to the server for a logical partition marked bootable, and of type 0x41 (PReP start). If a logical partition of this type does not exist, the logical partition IPL will fail.
*STMF (stream file)	This IPL type is used to start a logical partition from a kernel IBM i loaded in the IBM i integrated file system. Note that the integrated file system includes files on the optical (CD) drive on IBM i.

Unlinking virtual disk drives from an AIX logical partition

By unlinking virtual disk drives (network-server storage spaces) from an AIX logical partition that uses IBM i resources, you disconnect the virtual disk drives from the logical partition, making the virtual disk drives inaccessible to users. If you delete an AIX logical partition that uses IBM i resources, you must unlink all virtual disk drives from the logical partition before you delete the logical partition.

About this task

To unlink a virtual disk drive from an AIX logical partition that uses IBM i resources, follow these steps:

Procedure

1. Unlink disk drives from a logical partition by using IBM Navigator for i.

   If you prefer to use a character-based interface, go to step “2” on page 181.

   a) Vary off the NWSD for your logical partition.

   b) Click **Network** > **Windows Administration** > **Disk Drives**.

   c) Right-click the name of the disk drive that you want to unlink.

   d) Click **Remove Link**.

   e) Select a server from the list of linked servers.

   f) If you are unlinking a disk drive that you plan to relink later, clear **Compress link sequence**.

      You must relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.

   g) Click **Remove**.

   h) You have completed this procedure. Do not complete step “2” on page 181.

2. Unlink disk drives from a logical partition that uses a character-based interface:

   a) Vary off the NWSD for your logical partition.

   b) Type `RMVNWSSTGL` and press F4.

   c) In the **Network-server storage space** field, type the name of the storage space that you want to unlink and press Enter.

   d) In the **Network server description** field, type the name of the server from which you want to unlink the storage space and press Enter.

   e) If you are unlinking a linked disk drive that you plan to relink later, specify *NO in the **Renumber** field.

      Note: You must relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.

   f) Press Enter.

      Note: If you are uninstalling a logical partition, your next step is to delete the disk drive. For instructions, see “Deleting virtual disk drives for an AIX logical partition” on page 180. Otherwise, vary on the NWSD for your logical partition.

Saving AIX server objects in IBM i

When an AIX logical partition uses IBM i resources, IBM i stores AIX information in IBM i objects. IBM i can restore the objects correctly only if you save all objects for an AIX logical partition.

You can save these objects by using options of the IBM i GO SAVE command in the server.

- Option 21 saves the entire server.
- Option 22 saves server data, which includes objects in the QUSRSYS library.
- Option 23 saves all user data, which includes objects in the QFPNWSSTG library.

If you want to save a particular object, use the following table to see the location of that object on IBM i and the command to use.

Table 25. Objects to save for logical partitions with virtual disk

Object content	Object name	Object location	Object type	Save command
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network-server storage spaces in system auxiliary storage pool (ASP)	GO SAV, option 21 or 23
				SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network-server storage spaces in user ASP	SAV OBJ(('/QFPNWSSTG/stgspc') ('/dev/QASPnn /stgspc.UDFS')) DEV('/QSYS.LIB/TAP01.DEVD')

Table 26. Objects to save for all logical partitions with a server

Object content	Object name	Object location	Object type	Save command
Messages from the logical partition	Various	Various	Server message queue	GO SAVE, option 21 or 23
				SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
IBM i configuration objects for logical partitions	Various	QSYS	Device configuration objects	GO SAVE, option 21, 22, or 23
				SAVOBJ DEV (TAPO1)
Various	Various	QUSRSYS	Various	GO SAVE, option 21 or 23
				SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

Related information

Backup of the system image and user-defined volume groups

Installing system backups

Managing IBM i logical partitions that use i resources

You can manage IBM i logical partitions that uses i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Adding virtual disk units to an IBM i logical partition that uses i virtual I/O resources

You dynamically can add virtual disk units to an IBM i logical partition that uses i virtual I/O resources. This allows you to increase the storage capacity of the i logical partition that uses i virtual I/O resources when needed.

About this task

IBM i provides the ability to add virtual disks dynamically to another i logical partition. You can allocate disk space in the integrated file system and make it available to i without restarting the server or logical partition.

To add virtual disks dynamically to an IBM i logical partition that uses i virtual I/O resources, complete the following steps:

Procedure

1. Create a network-server storage space using the interface of your choice.

Interface	Actions
IBM Navigator for i	a. Expand **My Connections** > **your server** > **Network** > **Windows Administration**. b. Right-click the **Disk Drives** and select **New Disk**. c. In the **Disk drive name** field, specify the name that you want to give to the network-server storage space. d. In the **Description** field, specify a meaningful description for the network-server storage space. e. In the **Capacity** field, specify the size of the new network-server storage space in megabytes. To help you determine the size you want to use, see Installing, upgrading, or deleting IBM i and related software. f. Click **OK**.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides virtual I/O resources, type the command **`CRTNWSSTG`** and press F4. The Create NWS Storage Space (**`CRTNWSSTG`**) display opens. b. In the Network-server storage space field, specify the name you want to give to the network-server storage space. c. In the Size field, specify the size in megabytes for the new network-server storage space. d. In the Text description field, specify a meaningful description for the network-server storage space. e. Press Enter.

2. Add the new disk to the auxiliary storage pool (ASP) on the client IBM i logical partition.

   For instructions, see Adding disk units to an existing auxiliary storage pool (ASP).

3. Link the network-server storage space using the interface of your choice.

Interface	Actions
IBM Navigator for i	a. Expand **My Connections** > **your server** > **Network** > **Windows Administration**. b. Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**. c. Select the server to which you want to link the network-server storage space. d. Select one of the available data access types. e. Click **OK**.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides virtual I/O resources, type the command **`ADDNWSSTGL`** and press F4. The Add Network-Server Storage Link (**`ADDNWSSTGL`**) display opens. b. In the Network server description field, specify the name of the network server description (NWSD). c. In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without restarting the IBM i logical partition). d. In the Drive sequence number field, specify the link sequence position you want to use. e. Press Enter.

4. Activate the IBM i logical partition (if it is not already activated).

Linking a network-server storage space to a network server description

You can link a network-server storage space (NWSSTG) to a network server description (NWSD). This allows the NWSD and its associated logical partition to use the data stored on the NWSSTG.

About this task

You can link an NWSSTG to one NWSD. When you link an NWSSTG to an NWSD, you can set up the NWSD to have read-only access to the NWSSTG, or you can set up the NWSD to read or write to the NWSSTG.

To link an NWSSTG to an NWSD, follow these steps:

Procedure

1. At an IBM i command line on the IBM i logical partition that provides virtual I/O resources, type the command `ADDNWSSTGL` and press F4.
2. From the Add Server Storage Link display, provide the following information:

```
NWSSTG (Name)
NWSD (Name)
DYNAMIC (*YES)
DRVSEQNBR (*CALC)
```

Deleting network server descriptions for an IBM i logical partition that uses i virtual I/O resources

You can delete the IBM i network-server description (NWSD) for an i logical partition that uses i virtual I/O resources. When you delete the NWSD, all the configuration information for the i logical partition that uses i virtual I/O resources is deleted.

Before you begin

Before you start, ensure that you remove the disk from the auxiliary storage pool (ASP). For instructions, see Removing a disk unit from an auxiliary storage pool (ASP).

About this task

To delete the network-server description (NWSD) for an IBM i logical partition that uses i virtual I/O resources, follow these steps:

Procedure

1. On an IBM i control language (CL) command line on the i logical partition that provides virtual I/O resources, type the command WRKNWSD and press Enter.
2. Type 8 in the Opt field to the left of the Network Server and press Enter.
3. In the Work with Configuration Status display, if the status of the NWSD is not varied off, type 2 in the Opt field to the left of the Network Server and press Enter. Otherwise, go to the next step.
4. Press F3 to return to the previous display
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

Deleting virtual disk drives for an IBM i logical partition that uses i virtual I/O resources

You can delete a virtual disk drive from an IBM i logical partition that uses i virtual I/O resources. Deleting the virtual disk drive makes the space available once more to the i logical partition that provides the virtual disk resources. When you delete a virtual disk drive, all of the information on the virtual disk drive is erased.

Before you begin

Before you can delete a virtual disk drive, you must unlink the virtual disk drive from the network-server description (NWSD). For instructions, see “Unlinking virtual disk drives from an IBM i logical partition that uses i resources” on page 185.

About this task

To delete a virtual disk drive, follow these steps:

Procedure

Delete the disk drive using the interface that you prefer.

Interface	Actions
IBM Navigator for i	a. Click **Network** > **Windows Administration** > **Disk Drives**. b. Right-click the disk drive that you want to delete. c. Click **Delete** in the confirmation window.
IBM i character-based interface	a. At an IBM i control language (CL) command line, type `DLTNWSSTG` and press F4. b. Type the name of the disk drive in the Network-server storage space field and press Enter.

Using IPL types when running an IBM i logical partition that uses IBM i virtual I/O resources

The IPL source (IPLSRC) parameter on the network-server description (NWSD) determines the initial program that is loaded when the NWSD is varied on. For an IBM i logical partition that uses IBM i virtual I/O resources, the initial program is the load source. Ensure that the IPLSRC parameter specifies the location of the load source for the IBM i logical partition that uses IBM i virtual I/O resources.

You can set the IPLSRC parameter when you use the Create Network Server Description (CRTNWSD) command, and you can change the IPLSRC parameter when you use the Change Network Server Description (CHGNWSD) command.

The IPLSRC parameter has the following valid values.

IPLSRC values	Description
*Panel	The logical partition is started from the source indicated on the control panel.
*NWSSTG (network-server storage space)	This IPL type is used to start a logical partition from a virtual disk. The open firmware will find the load source in the virtual disk. If the firmware does not find the load source in the virtual disk, then the logical partition IPL will fail.
*STMF (stream file)	This IPL type is used to start a logical partition from a stream file loaded in the IBM i integrated file system of the IBM i logical partition that provides virtual I/O resources. Note that the integrated file system includes files on the optical (CD) drive on the IBM i logical partition that provides virtual I/O resources.

Unlinking virtual disk drives from an IBM i logical partition that uses i resources

By unlinking virtual disk drives (network-server storage spaces) from an IBM i logical partition that uses i resources, you disconnect the virtual disk drives from the logical partition, making the virtual disk drives inaccessible to users. If you delete an i logical partition that uses i resources, you must unlink all virtual disk drives from the logical partition before you delete the logical partition.

About this task

To unlink a virtual disk drive from an IBM i logical partition that uses i resources, follow these steps:

Procedure

Unlink disk drives from a logical partition using the interface that you prefer.

Interface	Actions
IBM Navigator for i	a. Vary off the NWSD for your logical partition. b. Click **Network** > **Windows Administration** > **Disk Drives**. c. Right-click the name of the disk drive that you want to unlink. d. Click **Remove Link**. e. Select a server from the list of linked servers. f. If you are unlinking a disk drive that you plan to relink later, clear **Compress link sequence**. You must relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence. g. Click **Remove**.
IBM i character-based interface	a. Vary off the NWSD for your logical partition. b. At an i command line on the i logical partition that provides virtual I/O resources, type `RMVNWSSTGL` and press F4. c. In the Network-server storage space field, type the name of the storage space that you want to unlink and press Enter. d. In the Network server description field, type the name of the server from which you want to unlink the storage space and press Enter. e. If you are unlinking a linked disk drive that you plan to relink later, specify *NO in the Renumber field. **Note:** You must relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the correct sequence. f. Press Enter. **Note:** If you are uninstalling a logical partition, your next step is to delete the disk drive. For instructions, see "Deleting virtual disk drives for an IBM i logical partition that uses i virtual I/O resources" on page 184. Otherwise, vary on the NWSD for your logical partition.

Saving IBM i server objects in i

When an IBM i logical partition uses i resources, the i logical partition that provides the virtual I/O resources stores the information for the i logical partition that uses the virtual I/O resources in i objects. i can restore the objects correctly only if you save all objects for an i logical partition that uses i virtual I/O resources.

You can save these objects by using options of the i GO SAVE command in the server.

- Option 21 saves the entire server.
- Option 22 saves system data, which includes objects in the QUSRSYS library.
- Option 23 saves all user data, which includes objects in the QFPNWSSTG library.

You can also use any save or restore command, or any Backup, Recovery, and Media Services (BRMS) function to save i server objects.

If you want to save a particular object, use the following table to see the location of that object on i and the command to use.

Table 27. Objects to save for logical partitions with virtual disk

Object content	Object name	Object location	Object type	Save command
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network-server storage spaces in system auxiliary storage pool (ASP)	GO SAV, option 21 or 23
				SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network-server storage spaces in user ASP	SAV OBJ(('/QFPNWSSTG/stgspc') ('/dev/QASPnn /stgspc.UDFS')) DEV('/QSYS.LIB/TAP01.DEVD')

Table 28. Objects to save for all logical partitions with a server

Object content	Object name	Object location	Object type	Save command
Messages from the logical partition	Various	Various	Server message queue	GO SAVE, option 21 or 23
				SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
IBM i configuration objects for logical partitions	Various	QSYS	Device configuration objects	GO SAVE, option 21, 22, or 23
				SAVOBJ DEV (TAPO1)
Various	Various	QUSRSYS	Various	GO SAVE, option 21 or 23
				SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

You can save and restore network server storage spaces on the i logical partition that uses i virtual resources; however, you cannot save and restore individual files.

Backing up and recovering IBM i logical partitions that use i virtual I/O resources
You can back up and recover an IBM i logical partition that uses resources from another i logical partition by using the GO SAVE operation.

Managing Linux logical partitions that use IBM i resources

You can manage Linux logical partitions that uses IBM i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Adding virtual disk units to a Linux logical partition
You can add virtual disk units dynamically to a Linux logical partition that uses IBM i resources. This allows you to increase the storage capacity of your AIX logical partition when needed.

About this task

Virtual disks simplify hardware configuration on the server because they do not require you to add additional physical devices to the server in order to run Linux. You can allocate up to 64 virtual disks to a Linux logical partition. Each virtual disk supports up to 1000 GB of storage. Each virtual disk appears to Linux as one actual disk unit. However, the associated space in the i integrated file system is distributed across the disks that belong to the i logical partition. Distributing storage across the disks provides the benefits of device parity protection through i. Therefore, you do not have to use additional processing resources and memory resources by setting up device parity protection through Linux.

IBM i provides the ability to dynamically add virtual disks to a Linux logical partition. You can allocate disk space in the integrated file system and make it available to Linux without restarting the server or logical partition. The Linux administrator can also configure the newly allocated disk space and make it available without restarting the server.

To add virtual disks dynamically to a Linux logical partition, complete the following steps:

Procedure

1. If you use IBM Navigator for i, create a network-server storage space using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration** .
   b) Right-click the **Disk Drives** and select **New Disk**.
   c) In the **Disk drive name** field, specify the name that you want to give to the network-server storage space.
   d) In the **Description** field, specify a meaningful description for the network-server storage space.
   e) In the **Capacity** field, specify the size of the new network-server storage space in megabytes.

      Refer to the installation documentation of your preferred Linux distributor to determine the size you want to use.
   f) Click **OK**.
   g) Continue with step "4" on page 188.
2. If you use a character-based interface, create a network-server storage space using the character-based interface:
   a) At an IBM i command line, type the command CRTNWSSTG and press F4.

      The Create NWS Storage Space (CRTNWSSTG) display opens.
   b) In the Network-server storage space field, specify the name you want to give to the network-server storage space.
   c) In the Size field, specify the size in megabytes for the new network-server storage space.

      Refer to the installation documentation of your preferred Linux distributor to determine the size you want to use.
   d) In the Text description field, specify a meaningful description for the network-server storage space.
   e) Press Enter.
3. If you use IBM Navigator for i, link the network-server storage space using IBM Navigator for i.
   a) Expand **My Connections** > **your server** > **Network** > **Windows Administration** .
   b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
   c) Select the server to which you want to link the network-server storage space.
   d) Select one of the available data access types.
   e) Click **OK**.
   f) Continue with step "5" on page 189.
4. If you use a character-based interface, link the network-server storage space using a character-based interface:
   a) At an IBM i command line, type the command ADDNWSSTGL and press F4.

      The Add Network-Server Storage Link (ADDNWSSTGL) display opens.
   b) In the Network server description field, specify the name of the network server description (NWSD).
   c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the Linux logical partition).
   d) In the Drive sequence number field, specify the link sequence position you want to use.

e) Press Enter.

5. If the Linux logical partition is not running, activate the Linux logical partition. Do not continue until the logical partition is running.
6. Log in to Linux using a user name with superuser (root) privileges.
7. Determine the host ID, SCSI bus, and logical unit number (LUN) for your new virtual disk drive.

   You can list the existing devices by typing the following command at the Linux command prompt: `cat /proc/scsi/scsi`. The following example shows sample output of the command:

   ```
   Attached devices:
   Host: scsi0 Channel: 00 Id: 00 Lun: 00
     Vendor: IBM      Model: VDASD NETSPACE   Rev: 0001
     Type:   Direct-Access                    ANSI SCSI revision: 04
   ```

   In this example, `NETSPACE` is the name of the network storage space for the displayed device. Look for the name of an existing network storage space on your Linux logical partition. Note the numeric part of the `Host:` value (host ID) and the `Channel:` (SCSI bus) and `Lun:` (logical unit number (LUN)) values for the existing network storage space. The new virtual disk drive will have the same host ID, SCSI bus, and LUN as the existing network storage space. For example, if the existing network storage space is as displayed in the preceding example output, then the new virtual disk drive will have a host ID of `0`, a SCSI bus of `0`, and a LUN of `0`.
8. Determine the SCSI ID for your new virtual disk drive.

   You can list the existing devices in table form by typing the following commands at the Linux command prompt:

   ```
   cd /proc/scsi/sg
   cat device_hdr; cat devices
   ```

   The following example shows sample output of the commands:

   ```
   host    chan    id      lun     type    opens   qdepth  busy    online
   0       0       0       0       0       2       30      0       1
   0       1       0       0       0       0       30      0       1
   ```

   Note the `host` (host ID), `chan` (SCSI bus), `id` (SCSI ID), and `lun` (logical unit number (LUN)) values for the existing devices. Find the devices that have the same host ID, SCSI bus, and LUN as the new virtual disk drive (as you determined in the previous step). Of those devices, find the device with the greatest SCSI ID. The new virtual disk drive will have a SCSI ID that is one greater than the greatest existing SCSI ID. For example, if the new virtual disk drive has a host ID of `0`, a SCSI bus of `0`, and a LUN of `0`, and the devices on your Linux logical partition are as listed in the example output above, then the new virtual disk drive will have a SCSI ID of `1`.
9. Type the following command at the Linux command prompt to add the virtual disk drive manually: `echo "scsi add-single-device host chan id lun" > /proc/scsi/scsi`.

   Use the following information to help you understand the arguments of the command:

   - `host` is the host ID.
   - `chan` is the SCSI bus.
   - `id` is the SCSI ID.
   - `lun` is the LUN.

   For example, if the new virtual disk drive is to have a host ID of `0`, a SCSI bus of `0`, a SCSI ID of `1`, and a LUN of `0`, you would type the command `echo "scsi add-single-device 0 0 1 0" > /proc/scsi/scsi` at the Linux command prompt.
10. At the Linux command prompt, type the following command to create a disk partition on the virtual disk drive: `fdisk /dev/sdb`.

    You must have superuser (root) privileges to run this command.

    The `Command (m for help):` prompt is displayed.
11. Type p at the prompt to see the current partition table for the virtual disk drive.

By default, the new virtual disk drive shows a single disk partition on the virtual disk.

For example,

```
Disk /dev/sdb: 64 heads, 32 sectors, 200 cylinders
Units = cylinders of 2048 * 512 bytes

Device Boot    Start       End    Blocks   Id   System
/dev/sdb1          1       199    203760    6   FAT16
```

12. Type d at the command prompt to delete the current partition and then create a new one.

    The default format for the disk partition is FAT16. Do not use a disk partition that is formatted as FAT16 on your virtual disk drive.

    The `Partition number (1-4):` prompt is displayed.
13. Type the disk partition number you want to delete and press Enter.
    In this example, you type a 1.

    The `fdisk` command indicates that the deletion is successful by displaying the command prompt.
14. Type n to create a new disk partition.

    The `Command action E extended P primary partition (1-4)` prompt is displayed.
15. Type p to create a primary disk partition on the virtual disk and press Enter.

    The `Partition number (1-4):` prompt is displayed.
16. Type 1 because this is the first partition on the virtual disk, and press Enter.

    The `First cylinder (1-200, default 1):` prompt is displayed.
17. Press Enter to use the default of 1 for the first disk cylinder.

    This uses the entire disk for this disk partition.

    The `Last cylinder or +size or +sizeM or +sizeK (1-200, default 200):` prompt is displayed.
18. Press Enter to use the default of 200 for the last disk cylinder.

    This uses the entire virtual disk for this partition.

    Note: The type of the partition defaults to Linux. If you need a different disk type (like Logical Volume Manager (LVM), or Linux Extended), type `t` to change the type of the partition.

    The `fdisk` command indicates that the partition creation is successful by returning the command prompt.
19. Type w to commit the changes to the disk structure and press Enter.

    The `fdisk` command writes the changes to the virtual disk drive. The `fdisk` command displays the following diagnostic message:

    ```
    The partition table has been altered!

    Calling ioctl() to re-read partition table.
    Syncing disks.
    ```

    After the operation is completed, the fdisk command returns the command prompt.
20. Format the disk partition using the Linux **mkfs** command .

    There are a number of optional parameters for the mkfs command, but typically the defaults satisfy most disk uses. To format the disk partition created in the previous steps, ensure that you are logged in with superuser (root) privileges and type the following command at a Linux command prompt:

    ```
    mkfs /dev/sdb1
    ```

    Since a single disk partition exists on the second virtual disk, the name of the disk is /dev/sdb1 (the sdb indicates that it is the second disk, and the 1 indicates that it is partition 1). The mkfs command displays the following diagnostic messages:

    ```
    mke2fs 1.28 (31-Aug-2002)
    Fileserver label=
    OS type: Linux Block size=1024 (log=0)
    ```

```
Fragment size=1024 (log=0)
51200 inodes, 204784 blocks
10239 blocks (5.00%) reserved for the super user
First data block=1
25 block groups
8192 blocks per group, 8192 fragments per group
2048 inodes per group
Superblock backups stored on blocks:
            8193, 24577, 40961, 57345, 73729

Writing inode tables: done
Writing superblocks and fileserver accounting information: done

This fileserver will be automatically checked every 29 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

21. Type the following command to create a directory that you can use to access the new file: `mkdir /mnt/data`
22. Type the following command to mount the virtual disk drive in the new directory: `mount /dev/sdb1 /mnt/data`
23. Add an entry to the /etc/fstab file using a Linux text editor, such as vi.
    For example, `/dev/sdb1 /mnt/data ext2 defaults 1 1`. This entry mounts the virtual disk every time you restart Linux.

Linking a network-server storage space to a network-server description
You can link a network-server storage space (NWSSTG) to one or more network-server descriptions (NWSDs). This allows the NWSDs and their associated logical partitions to use the data stored on the NWSSTG.

About this task

You can link an NWSSTG to an unlimited number of NWSDs. This is beneficial when multiple logical partitions need access to a single application.

When you link an NWSSTG to an NWSD, you can set up the NWSD to have read-only access to the NWSSTG, or you can set up the NWSD to read or write to the NWSSTG.

Attention: If more than one NWSD can write to the NWSSTG, ensure that only one NWSD can update the data at a time. Otherwise, changes made by one NWSD can be overwritten by another NWSD.

To link an NWSSTG to an NWSD, follow these steps:

Procedure

1. At an IBM i command line, type the command `ADDNWSSTGL` and press F4.
2. From the Add Server Storage Link display, provide the following information:

   ```
   NWSSTG (Name)
   NWSD (Name)
   DYNAMIC (*YES)
   DRVSEQNBR (*CALC)
   ```

3. Press F10 (Additional Parameters).
4. Enter the type of access the storage space will have.

Deleting network server descriptions for a Linux logical partition
You can delete the IBM i network server description (NWSD) for a Linux logical partition that uses i resources. When you delete the NWSD, all the configuration information for the Linux logical partition is deleted from i.

About this task

To delete the network-server description (NWSD) for a Linux logical partition, follow these steps:

Procedure

1. On an i control language (CL) command line, type the command WRKNWSD and press Enter.
2. Type 8 in the Opt field to the left of the Network Server and press Enter.
3. In the Work with Configuration Status display, if the status of the NWSD is not varied off, type 2 in the Opt field to the left of the Network Server and press Enter. Otherwise, go to the next step.
4. Press F3 to return to the previous display
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

Deleting virtual disk drives for a Linux logical partition
You can delete a virtual disk drive from a Linux logical partition that uses IBM i resources to make the space available to the i logical partition once more. When you delete a virtual disk drive, all of the information on the virtual disk drive is erased.

Before you begin

Before you can delete a disk drive, you must unlink it from the network-server description. For instructions, see "Unlinking virtual disk drives from a Linux logical partition" on page 193.

About this task

To delete a virtual disk drive, follow these steps:

Procedure

Delete the disk drive using the interface that you prefer.

Interface	Actions
IBM Navigator for i	Complete the following steps: a. Click **Network** > **Windows Administration** > **Disk Drives**. b. Right-click the disk drive that you want to delete. c. Click **Delete**. d. Click **Delete** in the confirmation window.
IBM i character-based interface	Complete the following steps: a. At an IBM i control language (CL) command line, type `DLTNWSSTG` and press F4. b. Type the name of the disk drive in the Network-server storage space field and press Enter.

Using IPL types when running Linux
The IPL source (IPLSRC) parameter on the network server description (NWSD) determines the initial program that is loaded when the NWSD is varied on. For a Linux logical partition that uses IBM i resources, the initial program is the kernel. Ensure that the IPLSRC parameter specifies the kernel location of the kernel for the Linux logical partition that uses i resources.

You can set the IPLSRC parameter when you use the Create Network Server Description (CRTNWSD) command, and you can change the IPLSRC parameter when you use the Change Network Server Description (CHGNWSD) command.

Note: The IPLSRC parameter also has the values A, B, and D, which are not valid for hardware that is used by IBM i logical partitions.

The IPLSRC parameter has the following valid values.

IPLSRC values	Description
*Panel	The logical partition is started from the source indicated on the control panel.
*NWSSTG (network-server storage space)	This IPL type is used to start a logical partition from a virtual disk. The open firmware will find the kernel in the virtual disk. The open firmware searches the first virtual disk connected to the server for a logical partition marked bootable, and of type 0x41 (PReP start). If a logical partition of this type does not exist, the logical partition IPL will fail.
*STMF (stream file)	This IPL type is used to start a logical partition from a kernel loaded in the IBM i integrated file system. Note that the integrated file system includes files on the optical (CD) drive on IBM i .

Unlinking virtual disk drives from a Linux logical partition

By unlinking virtual disk drives (network-server storage spaces) from a Linux logical partition that uses IBM i resources, you disconnect the virtual disk drives from the logical partition, making the virtual disk drives inaccessible to users. If you delete a Linux logical partition that uses i resources, you must unlink all virtual disk drives from the logical partition before you delete the logical partition.

About this task

To unlink a virtual disk drive from a Linux logical partition that uses i resources, follow these steps:

Procedure

1. Unlink disk drives from a logical partition using IBM Navigator for i.

   If you prefer to use a character-based interface, go to step "2" on page 193.

   a) Vary off the NWSD for your logical partition.

   b) Click **Network** > **Windows Administration** > **Disk Drives**.

   c) Right-click the name of the disk drive that you want to unlink.

   d) Click **Remove Link**.

   e) Select a server from the list of linked servers.

   f) If you are unlinking a disk drive that you plan to relink later, uncheck **Compress link sequence**.

      You must relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.

   g) Click **Remove**.

   h) You have completed this procedure. Do not complete step "2" on page 193.

2. Unlink disk drives from a logical partition using a character-based interface:

   a) Vary off the NWSD for your logical partition.

   b) Type `RMVNWSSTGL` and press F4.

   c) In the Network-server storage space field, type the name of the storage space that you want to unlink and press Enter.

   d) In the Network server description field, type the name of the server from which you want to unlink the storage space and press Enter.

   e) If you are unlinking a linked disk drive that you plan to relink later, specify *NO in the Renumber field.

      Note: You must relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.

f) Press Enter.

Note: If you are uninstalling a logical partition, your next step is to delete the disk drive. For instructions, “Deleting virtual disk drives for a Linux logical partition” on page 192. Otherwise, vary on the NWSD for your logical partition.

Saving Linux server objects in IBM i
When a Linux logical partition uses IBM i resources, i stores Linux information in i objects. IBM i can restore the objects correctly only if you save all objects for a Linux logical partition.

You can save these objects by using options of the i GO SAVE command in the server.

- Option 21 saves the entire server.
- Option 22 saves server data, which includes objects in the QUSRSYS library.
- Option 23 saves all user data, which includes objects in the QFPNWSSTG library.

If you want to save a particular object, use the following table to see the location of that object on i and the command to use.

Table 29. Objects to save for logical partitions with virtual disk

Object content	Object name	Object location	Object type	Save command
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network-server storage spaces in system auxiliary storage pool (ASP)	GO SAV, option 21 or 23
				SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network-server storage spaces in user ASP	SAV OBJ(('/QFPNWSSTG/stgspc') ('/dev/QASPnn /stgspc.UDFS')) DEV('/QSYS.LIB/TAP01.DEVD')

Table 30. Objects to save for all logical partitions with a server

Object content	Object name	Object location	Object type	Save command
Messages from the logical partition	Various	Various	Server message queue	GO SAVE, option 21 or 23
				SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
IBM i configuration objects for logical partitions	Various	QSYS	Device configuration objects	GO SAVE, option 21, 22, or 23
				SAVOBJ DEV (TAPO1)
Various	Various	QUSRSYS	Various	GO SAVE, option 21 or 23
				SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

Backing up and recovering Linux logical partitions that use IBM i virtual I/O resources
When you create a Linux logical partition that uses resources from an IBM i logical partition, you can manage backup and recovery using IBM i control language (CL) commands, Linux commands, or a combination of the two.

To save Linux data in a logical partition that uses i resources to a shared tape drive and restore the data from the tape drive, you can use either the Linux **tar** command or the i Save (SAV) and Restore (RST) commands. You can also use the **tar** command to save your data to a file. If you use the **tar** command to save data, the only way you can restore that data is by using the **tar** command again. Similarly, if you use the SAV command to save data, the only way you can restore that data is by using the RST command. The two methods of backing up and restoring data are not compatible.

The following restrictions apply:

- To use the tape device from Linux, you must vary the tape off under IBM i.
- Saving the storage space is typically faster than saving by using the **tar** command, but it does not provide file-level backup and recovery.
- Linux does not support switching tapes in a library device. You can only use the tape that is currently in the device.
- You cannot save IBM i data and **tar** data on the same tape volume.

Backing up and recovering files using the tar command
The most common data backup utility in Linux is the **tar** (tape archive) utility. Use the Linux **tar** command if you have Linux installed on a dedicated disk or if you cannot vary off a Linux logical partition while you are backing up data.

Backups using the Linux **tar** command are at the file level. They save only the files and directories that the **tar** command specifies. Therefore, you cannot use the **tar** command to save Linux data that is not in the file server. For example, you cannot save a kernel in the PowerPC® Reference Platform (PReP) start logical partition by using the **tar** command.

One advantage of the **tar** command is that it supports incremental backups and backup of special devices, which is not common for `tar` implementations. Also, the **tar** command backs up files without regard to the underlying file system type.

Saving to and restoring from a tape device
Use these procedures to save and restore Linux files between a Linux logical partition that uses IBM i resources and a shared tape drive.

Before you begin

Ensure that your Linux data is in the file server.

Linux typically treats tape as a *character device* that it can quickly read from or write to in long streams of data, but cannot quickly access to find specific data. By contrast, Linux treats a disk or CD as a *block device* that it can read from or write to quickly at any point on the device, making it suitable for the **mount** command.

About this task

Complete the following steps to save and restore Linux files between a logical partition that uses IBM i resources and a shared tape drive:

Procedure

1. Type the following command: `tar -b 40 -c -f /dev/st0` *`files`*

   Use the following descriptions to help you understand the arguments of this command:

   - `tar` is the command name (the contraction of "tape archive").

- `-b 40` is the block size in sectors. This argument specifies that Linux is to write the archive stream in blocks of 40 sectors (20 KB). If you do not specify a value for this argument, the default value is 20 sectors (10 KB), which does not perform as well over virtual tape as does a value of 40.
- `-c` is the command action to create. This argument specifies that the **tar** command creates a new archive or overwrites an old one (as opposed to restoring files from an archive or adding individual files to an existing archive).
- `-f /dev/st0` is the virtual tape device and number. This argument specifies that the command uses virtual tape 0 on the server. After the **tar** command runs, the tape device is closed and the tape is rewound. To save more than one archive on the tape, you must keep the tape from rewinding after each use, and you must position the tape to the next file marker. To do this, specify the *`nst0`* (nonrewinding virtual tape) device instead of *`st0`*.
- *`files`* are the names of the files and directories that you plan to save.

You have now saved Linux data from a logical partition that uses IBM i resources to the shared tape drive.

2. Type the following command: `tar -b 40 -x -f /dev/st0` *`files`*

The **-x** (extract) argument replaces the **-c** (create) argument in the **tar** command used in step "1" on page 195.

You have now restored Linux data from the shared tape drive to a logical partition that is sharing resources.

Saving to and restoring from a file
You can save and restore Linux files between a Linux logical partition that uses IBM i resources and a tar file.

Saving to a file

The following is an example of using the **tar** command to save to a file.

```
tar -cvf /tmp/etc.tar /etc
```

Use the following descriptions to help you understand the arguments of this command:

tar
The command name.

c
Create a tar file.

v
Verbose. This argument shows the files that are being added to the tar file.

f
The data immediately following `f` is the name of the tar file.

/tmp/etc.tar
The name of the tar file.

/etc
An object to be added to the tar file. Because /etc is a directory, the utility adds all the contents of the directory and its subdirectories to the tar file.

After you create the tar file, you can save it to an offline medium in several ways. For example, you can save the tar file to a virtual tape device or a directly attached tape device. You can also copy the tar file to the integrated file system and save it at a later time.

You can save the data on a Linux logical partition to a tar file during normal server usage. You can automate and start the **tar** utility by using the **cron** (chronology) daemon on the logical partition. The **cron** daemon is a scheduling mechanism for Linux. You can also use the **tar** utility to schedule a single backup request. For example, if you want to use the tar utility to back up the /etc directory at 10 p.m. on 19 September, you can type the following command:`at 10pm Sep 19 -f tar.command`.

Restoring from a file

The following is an example of using the **`tar`** command to restore from file: `tar -xvf /tmp/etc.tar /etc`. The **-x** (extract) argument replaces the **-c** (create) argument in the **`tar`** command used to save the files.

Backing up and recovering Linux logical partitions using i commands
If you have a Linux logical partition that uses IBM i resources, tools are available in i for backup and recovery. You can use the Save (SAV) and Restore (RST) control language (CL) commands to save and restore entire virtual disks in their current state.

The SAV command saves the directory that has the same name as the virtual disk under the QFPNWSSTG directory in the integrated file system. This method of backup and recovery is most effective if the Linux kernel is saved in a PowerPC Reference Platform (PReP) start logical partition on the virtual disk. On most Linux distributions, this usually occurs as part of a default installation.

Backups of storage spaces using i commands are at drive level. This means that i backs up the entire contents of a virtual disk, or network storage space, rather than individual files. Thus, the correct SAV command backs up any information on the drive, including a kernel in the PReP start logical partition.

If you save the Linux kernel in a PReP logical partition, you can restore and start the logical partition after a total system re installation. You can also transport and restore saved virtual disks to other servers using File Transfer Protocol (FTP) and tape.

Save Linux data by using IBM i SAV
You can save data for a Linux logical partition that uses IBM i resources by using the Save (SAV) i CL command.

About this task

On IBM i, your data is in a network-server storage space.

Saving and restoring individual Linux files using IBM i commands requires that you use the QNTC directory in the integrated file system. You can access the files that you save and restore by using a file share. You can define the fiile share by using Samba on Linux, and you can access the file share by using QNTC.

To save data for a Linux logical partition that uses IBM i resources by using the Save (SAV) IBM i CL command, follow these steps:

Procedure

1. At the IBM i command line, enter the Save (SAV) command.
2. On the Save display, enter the following parameter values:
   a) In the **Device** field, enter the associated IBM i device description.
      To save to a file in a library like QGPL, enter `/qsys.lib/qgpl.lib/myfile.file`. For example, if your tape device is named TAP01, enter `/qsys.lib/tap01.devd`.
   b) In the **Objects: Name** field, enter the server, share, or file.
      For example, if your server is named MYSERVER, your share is named MYSHARE, and it contains all of the directories and files that need to be saved, enter `/QNTC/MYSERVER/MYSHARE`.
3. At the IBM i command line, enter the Display Save File (DSPSAVF) command to verify that the changed save file exists.
4. In the Option field by the new save file name, enter 5 (Display) to display a list of the stream files in the save file.

Restore Linux data using i RST
You can restore data for a Linux logical partition that uses i resources by using the Restore (RST) IBM i CL command.

About this task

Saving and restoring individual Linux files using IBM i commands requires that you use the QNTC directory in the integrated file system. You can access the files that you save and restore by using a file share. You can define the file share by using Samba on Linux, and you can access the file share by using QNTC.

Restore (RST) is the i command to restore Linux files from the shared tape drive of the logical partition that shares resources. On the Restore Object display, enter the following parameter values:

Procedure

1. To restore from a tape device, enter the associated i device description in the **Device** field.
   For example, if your tape device is named TAP01, enter `/qsys.lib/tap01.devd`.
2. To restore from a save file in library QGPL, enter the associated file name.
   For example, `/qsys.lib/qgpl.lib/myfile.file`.
3. In the **Objects: Name** field, enter the server, share, or file.
   For example, if your server is named MYSERVER, your share is named MYSHARE, and it contains all of the directories and files that need to be restored, enter `/QNTC/MYSERVER/MYSHARE`.

Backing up the network server description and virtual disk drives associated with a Linux logical partition
Learn about how to back up the data for a Linux logical partition that uses IBM i resources.

Backing up the data for a Linux logical partition that uses IBM i resources is different from backing up the data for a Linux logical partition that uses its own resources. When you install the logical partitions with virtual disk, the IBM i logical partition that shares resources creates a network server description and creates disk drives for your Linux logical partition that you need to back up. Some of the disk drives are server-related (the installation and server drives), while others are user-related. Because your Linux logical partition might consider the disk drives to be a unified server, you must save all the disk drives and the network server description so they restore correctly.

With the implementation of a logical partition, you can save and restore virtual disks as IBM i network-server storage space objects. These objects are saved as part of the server when you perform a full server backup. You can also specifically save the network server description and storage spaces that are associated with a logical partition on a server. Daily backup of the server drive is a good practice.

Building a rescue image on a network storage space
You can build a rescue image on a network storage space (NWSSTG) to assist you in checking and repairing a faulty Linux installation.

Before you begin

A *rescue image* is a disk image that contains the Linux kernel, a shell, and the diagnostic tools, drivers, and other utilities that would be useful for checking and repairing a faulty Linux installation. Many Linux distributors include a rescue image on their installation disks. One rescue solution for a logical partition is to create a small NWSSTG that can remain on the integrated file system solely for the purpose of rescuing logical partitions. You can install a rescue image to the NWSSTG when you create your logical partition.

Before creating a rescue image on network storage, it is important to document the configuration information for each of your logical partitions.

1. Document the drive configuration information, which is located in the /etc/fstab file.
2. Capture the networking information that is reported when you run the **ifconfig** command.
3. Create a list of the modules that are needed by each logical partition. You can see which modules are in use by using the **lsmod** command from within Linux. Use the information obtained from the

commands and files listed above to determine which files to store on your rescue network storage space.

About this task

To build a rescue image on an NWSSTG, follow these steps:

Procedure

1. Determine how much network storage space you need to build the rescue image.

   Consult your Linux documentation to see how much space is required for a minimum installation of your distribution, and add enough space to create a swap partition (a PowerPC Reference Platform (PReP) start partition) and to install any extra software that you would like to have available in your rescue image. For example, if the documentation states that a minimum server installation is 291 MB, create a storage space of 425 MB.
2. Create a network storage space (CRTNWSSTG) of the size you determined for the rescue image.

   You might want to make a note in the storage space description field that indicates which distribution was used to make the rescue image and warns that it should be saved.
3. Link this storage space to a network server description (NWSD).

   You do not need to create a new NWSD for this step. You could unlink an existing storage space and temporarily link your rescue storage space to any of your existing NWSDs.
4. Start the installation server for your distribution as described in the documentation and follow the prompts.

   To partition your installation manually, ensure that you create a PReP start partition. At the point where you select the packages to install, select the minimum number of packages supported. The name for the package group varies by distribution.
5. Allow the installer to complete its package installation and configuration.

   After installation has finished, the installer starts the rescue image for you.
6. Verify that the rescue image has all the utilities that you need.

   For a logical partition, at a Linux command prompt, type `rpm -qa | grep ibmsis` to make sure that the utilities that work with the integrated disk are available.
7. Ensure that the device drivers that your logical partitions require are installed.
   For example, verify that pcnet32 is installed for Ethernet devices, or that olympic is installed for token-ring devices. The kernel modules that have been compiled can be found in the /lib/modules/ kernel version/kernel/drivers directory or in directories under that directory.
8. Install any other special drivers or software packages that your logical partitions require.
9. Use File Transfer Protocol (FTP) to send the files with the configuration information for your other logical partitions to the rescue server network storage space.
10. Install the kernel manually (if you are required to do so by your Linux distribution).

    For details regarding installing the kernel, consult the appropriate installation documentation for your distribution.
11. Make note of the path to the root partition on the rescue-storage space.

    You must use this information to start the rescue network storage space from the network. To determine the root partition, type the command `cat /etc/fstab`. The partition that has a forward slash (/) in the second column is your root partition. For further assistance in determining the root partition, see the documentation for your distribution.

What to do next

You can shut down your logical partition by typing `shutdown -h` now and varying off the logical partition after the shutdown has completed. After the logical partition has varied off, you can unlink the rescue storage space and relink the normal storage space for the NWSD.

Using a rescue image from a network-server storage space

You can use a Linux rescue image on a network-server storage space (NWSSTG) to repair a Linux logical partition that uses IBM i resources. A *rescue image* is a disk image that contains the Linux kernel, a shell, and the diagnostic tools, drivers, and other utilities that would be useful for checking and repairing a faulty Linux installation.

About this task

To use the rescue image that you built on the NWSSTG, use the following steps:

Procedure

1. Disconnect the virtual storage space for the failed logical partition (if applicable) by using the Work with NWS Storage Spaces (WRKNWSSTG) command.
2. Connect your rescue storage space as the first drive to the network server description (NWSD), and reconnect the original storage space (where applicable) as the second drive.
3. Edit the NWSD for the failed partition so that it starts from IPL source *NWSSTG. Also, edit the IPL Parameters field to reflect the root partition on the rescue storage space. For most distributions, this is a parameter such as `root=/dev/sda3` or `root=/dev/vda1`. For assistance, see the documentation for your Linux distribution.
4. Restart the partition.
5. If the existing root partition is on a dedicated disk, you might need to insert the `ibmsis` driver using the `insmod ibmsis` command.
6. Create a mount point to which you will mount the root partition of the network storage space that you are trying to rescue. You can use a command such as `mkdir /mnt/rescue`.
7. Mount the root partition of the network storage space that you are trying to rescue. Mount a drive using the command `mount -t` ***`partition-type partition-location mount-point`***, where the partition type is the format of the partition such as ext2 or reiserfs, the partition location is similar to /dev/sdb3 (for non-devfs disk partitions), /dev/sd/disc1/part3 (for devfs disk partitions), or /dev/sda2 (for a partition on a dedicated disk).
8. The drive that you are trying to rescue, when using virtual disk, will be the second drive rather than the first drive. (That is, if the drive was /dev/sda3 when the partition was running normally, it will be /dev/sdb3 in the rescue server.)
9. Use the documentation or the configuration files you created when you created the rescue NWSSTG to help you determine the device for the root of the partition you are trying to rescue. Your mount point will be similar to /mnt/rescue if you use the previous example.

What to do next

You can either use the rescue tools provided in your rescue storage space against the mount point you have created or you can work on the partition that you are rescuing from within its own storage space. If rescuing the image from its own storage space, change the root directory for that partition using the `chroot` *`mount-point`* command.

Backing up network server descriptions for a Linux logical partition

When you save the storage space objects that are associated with a logical partition that uses virtual disks, you must also save the network server description (NWSD). Otherwise, a logical partition might not be able to re-establish items such as the file-system permissions for the logical partition.

About this task

Use the Save Configuration (SAVCFG) command to save the network server description:

Procedure

1. On the IBM i command line, type SAVCFG.
2. Press Enter to save the NWSD configuration.

What to do next

The Save Configuration command (SAVCFG) saves the objects associated with an NWSD, including the line descriptions and network-server storage space link information. SAVCFG does not save the storage spaces associated with this server. You can use the Save Object (SAV) command to save the storage spaces.

Restoring network-server descriptions for a Linux logical partition
In a disaster-recovery situation, you would restore all the configuration objects, which include the network-server description (NWSD) for your logical partition. In some situations, you must specifically restore the NWSD. For example, you must restore the NWSD when you migrate to new hardware.

Before you begin

To have IBM i automatically relink disk drives within the integrated file system to the restored NWSD, restore those disk drives first.

About this task

To restore the NWSD, use the Restore Configuration (RSTCFG) command:

Procedure

1. On an IBM i command line, type RSTCFG and press F4 (Prompt).
2. In the **Objects** field, specify the name of the NWSD.
3. In the **Device** field, specify which device you are using to restore the NWSD.

   If you are restoring from media, specify the device name. If you are restoring from a save file, specify *SAVF and identify the name and library for the save file in the appropriate fields.
4. Press Enter to restore the NWSD.
5. When you have restored the NWSD and all of its associated storage spaces, start (vary on) the logical partition.

Synchronizing the hypervisor and Service Processor time-of-day clocks to Time Reference Partition

You can ensure that the time-of-day clocks used by the hypervisor and Service Processor are accurate through the usage of one or more Time Reference Partitions (TRP). Whenever the time-of-day changes in the TRP, the hypervisor and the service processors time-of-day are updated to match the time that is specified by the TRP. The time of the TRP can be changed manually or can be managed by using the Network Time Protocol (NTP) support. NTP can be used to automatically ensure consistent time-of-day across multiple servers. When you designate a logical partition as a TRP, you must choose a logical partition that cannot be migrated to another server such as a Virtual I/O Server (VIOS) partition. More than one TRP can be specified per server and the longest running TRP is recognized as the TRP of the system.

About this task

The service processor uses its time-of-day clock for time stamps of various error logs and events. The service processor has a battery backed clock so if an electrical outage occurs, the server can maintain the current time-of-day. Whenever a server is powered on, the hypervisor time-of-day clock is initialized from the time-of-day clock of the service processor. The hypervisor uses its time-of-day clock whenever a new partition is created. Newly created partitions start with the time-of-day clock of the hypervisor. After a logical partition is created, the time-of-day clock of the logical partition is separate from the time-of-day clock of the hypervisor (changes to the time of the hypervisor does not affect existing partitions).

To enable TRP on a partition, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select logical partition and click **Actions** > **View Partition Properties**. The **Properties** page is displayed.
4. Click the **General** tab.
5. Click the **Advanced** tab. In the **Advanced Settings** area, select the **Enable Time Reference** check box.
6. Repeat steps 3 through 5 for any additional logical partitions.

Enabling user mode access to the hardware accelerator

You can enable user mode access to the hardware accelerator by using the Hardware Management Console (HMC). The HMC must be at Version 9.1.940, or later.

About this task

The GNU zip (gzip) accelerators are compression and decompression cards that are used to increase server performance and network I/O efficiency. Quality of service (QoS) credits is a mechanism that is used to give logical partitions access to shared hardware accelerators. You can verify whether the server supports enabling user mode access to the hardware accelerator by using the **`lssyscfg`** command.
To view the supported hardware accelerators types, along with the corresponding maximum Hardware Accelerator QoS, currently available Hardware Accelerator QoS, and pending available Hardware Accelerator QoS credits of the server, you can use the **`lshwres`** command. You can also use the **`lshwres`** command to view the amount of gzip QoS that is assigned to a logical partition. You can enable the gzip quality of service (QoS) credits for an logical partition AIX, Linux, or Virtual I/O Server logical partition that are either in the `Activated` or `Not Activated` states by using the **`chhwres`** command. You cannot enable the QoS credits for an IBM i logical partition.

The logical partition is assigned credits only when the following conditions are met:

- The server supports enabling the Hardware Accelerator.
- The server supports the Hardware Accelerator type that you specify.
- The operating system supports enabling QoS credits dynamically.
- Sufficient amount of Hardware Accelerator QoS credits are available that can be assigned to the logical partition.

To remote restart a logical partition that is enabled with gzip QoS credits, the destination server must support user mode access to the hardware accelerator. Additionally, the following conditions apply:

- When the HMC that manages the destination server is at version 9.1.0, or earlier, the remote restart operation succeeds with the logical partition losing the gzip QoS credits after the remote restart operation completes.
- The remote restart operation succeeds when there is a sufficient amount of available gzip QoS credits on the destination server and the HMC that manages the destination server is at version 9.1.940, or later. The logical partition is restarted with the same amount of gzip QoS credits that was available at the source server.
- When the destination server has an insufficient amount of available gzip QoS credits, and the server is managed by an HMC at version 9.1.940, or later, the remote restart operation succeeds and the logical partition is assigned partial gzip QoS credits that depends on the amount of available gzip QoS credits on the destination server.

Related information

chhwres command

lshwres command

lssyscfg command

Performance considerations for logical partitions

You can manage and enhance the performance of logical partitions so that your system uses its resources in the most efficient manner.

You can manage and enhance the performance of a AIX logical partition by configuring the AIX operating system.

Managing IBM i performance ensures that your managed system is efficiently using resources and that your managed system provides the best possible services to you and to your business. Moreover, effective performance management can help you quickly respond to changes in your managed system and can save on expenses by postponing costly upgrades and service fees.

Related information

AIX Performance management

Performance Tools Guide and Reference

Performance Toolbox Version 2 and 3 Guide and Reference

Power Systems Capacity on Demand

Adjusting the Active Memory Expansion configuration to improve performance

You can run the Active Memory Expansion planning tool to generate performance statistics for an AIX logical partition that uses Active Memory Expansion. Then, you can change the Active Memory Expansion factor, the memory assignment, or the processor assignment of the logical partition to improve its performance.

About this task

To adjust the Active Memory Expansion configuration to improve performance, complete the following steps:

Procedure

1. Run the Active Memory Expansion planning tool, which is the **amepat** command, from the AIX command line interface.

   When you run the planning tool on a workload that currently uses Active Memory Expansion, the tool generates a report that provides the following information:

   - Various statistics about memory compression and processor consumption.
   - Several alternative configuration possibilities for Active Memory Expansion on the logical partition.
   - Recommended configuration to improve the performance of Active Memory Expansion on the logical partition.

   Tip: You can view more detailed statistics about memory compression and processor consumption by using the **vmstat**, **lparstat**, **svmon**, and **topas** commands.
2. Perform one or more of the following tasks to adjust the configuration:
   - Dynamically change the Active Memory Expansion factor that is set for the logical partition. For instructions, see “Changing the Active Memory Expansion factor for AIX logical partitions” on page 150.
   - Dynamically add, move, or remove memory to or from the logical partition. For instructions, see one of the following tasks:
     - For logical partitions that use dedicated memory, see “Managing dedicated memory dynamically” on page 149.

  - For logical partitions that use shared memory, see "Adding and removing logical memory dynamically to and from a shared memory partition" on page 151.
- Dynamically add, move, or remove processor resources to or from the logical partition. For instructions, see "Managing processor resources dynamically" on page 153.

Performance considerations for shared memory partitions

You can learn about performance factors (such as shared memory overcommitment) that influence the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). You can also use shared memory statistics to help you determine how to adjust the configuration of a shared memory partition to improve its performance.

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

A shared memory configuration is considered over committed when the sum of the logical memory that is assigned to all of the shared memory partitions is greater than the amount of physical memory in the shared memory pool.

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

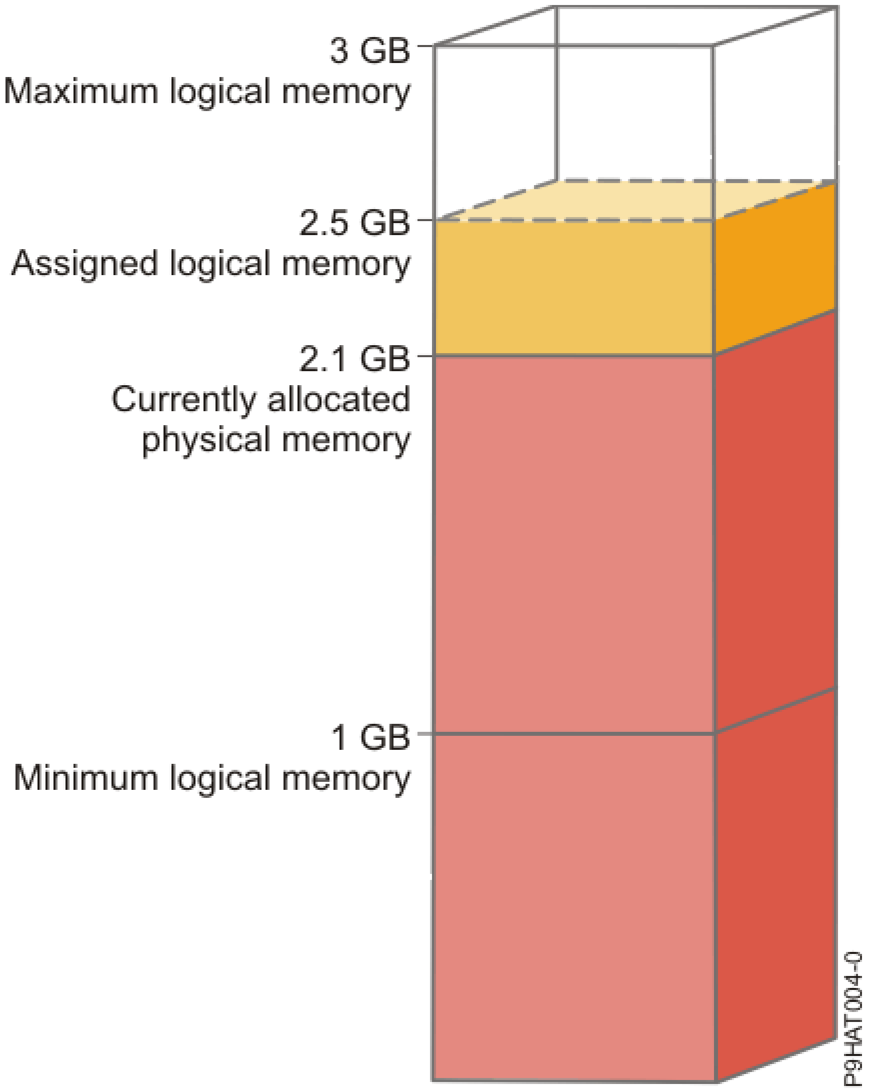

Figure 7. A shared memory partition in a logically overcommitted memory configuration

The figure shows a shared memory partition that is assigned 2.5 GB of logical memory. Its maximum logical memory is 3 GB and its minimum logical memory is 1 GB. The figure also shows that the amount of physical memory that is currently allocated to the shared memory partition from the shared memory pool is 2.1 GB. If the workload that runs in the shared memory partition currently uses 2.1 GB of memory and requires an additional 0.2 GB of memory, and the shared memory pool is logically overcommitted, the hypervisor allocates an additional 0.2 GB of physical memory to the shared memory partition by assigning memory pages that are not currently in use by other shared memory partitions.

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

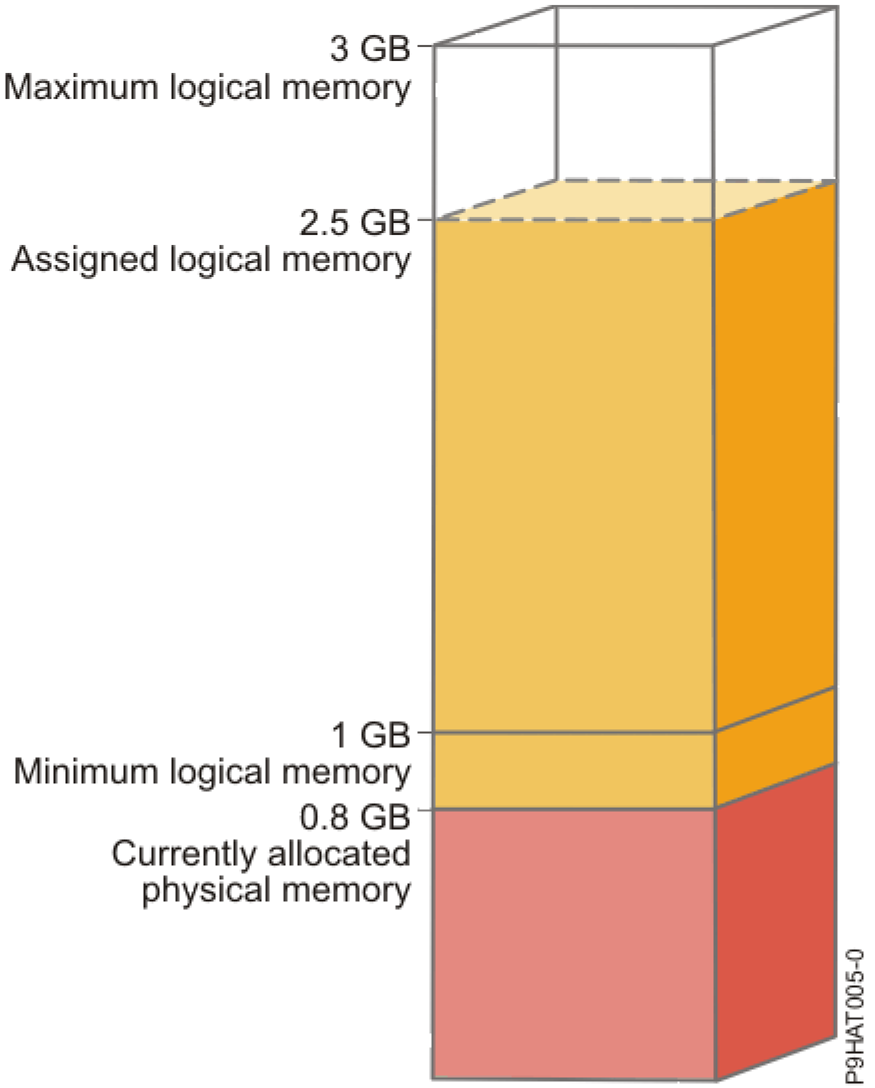

Figure 8. A shared memory partition in a physically over committed memory configuration

The figure shows a shared memory partition that is currently allocated 0.8 GB of physical memory and assigned 2.5 GB of logical memory. If the workload that runs in the shared memory partition currently uses 0.8 GB of memory and requires an additional 1.5 GB of memory, and the shared memory pool is physically over committed, the hypervisor stores 1.5 GB of the shared memory partition's memory in its paging space device.

When the shared memory partition needs to access data on the paging space device, the hypervisor directs a paging VIOS partition to read the data from the paging space device and write the data to the shared memory pool. The more memory that the hypervisor must store on the paging space device, the more often the hypervisor and paging VIOS partition need to read and write data between the paging space device and the shared memory pool. Compared to directly accessing data that is stored in the shared memory pool, it takes more time to access data that is stored in the paging space device. Thus, in general, the less over committed the memory configuration of a shared memory partition, the better its performance.

The operating systems that run in the shared memory partitions help improve the performance of the shared memory partitions with over committed memory configurations by providing the hypervisor with information about how the operating system uses the physical memory that is allocated to it. Using this information, the hypervisor can store data that the operating system accesses the least often in the paging space device and store the data that the operating system accesses the most often in the shared memory pool. This reduces the frequency that the hypervisor needs to access the paging space device and increases the performance of the shared memory partition.

Related concepts

Factors that influence the performance of shared memory partitions
In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Example: A shared memory configuration that is logically overcommitted
When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

Example: A shared memory configuration that is physically overcommitted
When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

Shared memory distribution
The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Related reference

Performance statistics for shared memory
The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

The following table describes the types of workloads that are appropriate to run in shared memory configurations that are logically and physically over committed. It also describes the types of workloads that are not appropriate to run in a shared memory configuration.

Table 31. Workloads to run in logically over committed configurations, physically over committed configurations, and dedicated memory configurations

Workloads for logically overcommitted configurations	Workloads for physically overcommitted configurations	Workloads for dedicated memory configurations
• Workloads that peak at opposite and varying times. • Workloads with memory residency requirements that have a low average. • Workloads that do not have a sustained load. • Logical partitions that serve as failover and backup logical partitions when configured on the same server as their primary counterparts. • Test and development environments.	• Workloads that run the AIX operating system and use the file cache. • Print servers, file servers, network applications, and other workloads that are less sensitive to I/O latency. • Workloads that are inactive most of the time.	• Workloads with high quality of service criteria. • Workloads that consistently use memory resources due to sustained peak load. • High-performance computing (HPC) workloads.

In addition to the degree to which the memory configuration of a shared memory partition is over committed, the following factors can influence the performance of a shared memory partition:

- The workload that runs in a shared memory partition, the number of virtual adapters that are assigned to the shared memory partition, and the I/O entitled memory set for the shared memory partition all directly affect the performance of I/O devices. These factors can cause I/O devices to operate at their minimum memory requirements rather than their optimal memory requirements. This can cause delays in I/O operations.
- The amount of I/O entitled memory that is required for optimal performance depends on the workload and number of adapters configured.
- The operating systems that run in shared memory partitions cannot use memory affinity. Some applications rely on memory affinity to improve their performance.
- The shared memory partition might be suspended if it attempts to access data on its paging space device when the following situations occur simultaneously:
  - The paging VIOS partition becomes unavailable. For example, you shut down the paging VIOS partition or the paging VIOS partition fails.
  - The shared memory partition is not configured to use redundant paging VIOS partitions to access its paging space device.

Related concepts

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Related reference

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Where to view statistics	Statistics to view
HMC utilization data	• Statistics about the shared memory pool, such as: – Size of the shared memory pool – Total amount of memory that is over committed – Total amount of logical memory that is assigned to the shared memory partitions – Total amount of I/O entitled memory that is assigned to the shared memory partitions – Total amount of physical memory that the shared memory partitions currently use for their I/O devices – Amount of memory from the shared memory pool that the hypervisor uses to manage the shared memory partitions – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the shared memory partitions, such as: – Amount of logical memory assigned to the shared memory partition – Amount of physical memory from the shared memory pool that is allocated to the shared memory partition – Amount of memory that is over committed – I/O entitled memory assigned to the shared memory partition – Amount of physical memory that the shared memory partition currently uses for its I/O devices – Memory weight of the shared memory partition

Where to view statistics	Statistics to view
IBM i See IBM® i to view shared memory statistics in IBM i.	• Statistics about the shared memory pool, such as: – Total number of page faults for all of the shared memory partitions – Total time, in milliseconds, that the processors waited for page faults to be resolved – Total physical memory, in bytes, that is assigned to the shared memory pool – Sum of the logical memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the I/O entitled memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the physical memory, in bytes, that the shared memory partitions that are active currently use for their I/O devices • Statistics about the shared memory partition, such as: – Memory weight of the shared memory partition – Amount of physical memory, in bytes, from the shared memory pool that is currently used by the shared memory partition – Number of times that the shared memory partition waited for a page fault – The time, in milliseconds, that the shared memory partition waited for page faults to be resolved – Maximum amount of memory, in bytes, that the shared memory partition can assign to data areas that are shared between the operating system and the server firmware – I/O entitled memory assigned to the shared memory partition – Minimum amount of physical memory, in bytes, required for all of the configured I/O devices to operate – Optimal amount of physical memory, in bytes, required for I/O devices to maximize throughput performance – Amount of physical memory, in bytes, that the shared memory partition currently uses for its I/O devices – Highest amount of physical memory, in bytes, that the shared memory partition has used for its I/O devices since the last time the shared memory partition was activated or since the last time the memory statistics were reset, whichever is most recent – Number of delayed I/O operations since the last time the shared memory partition was activated

Where to view statistics	Statistics to view
Linux View memory statistics for Linux in the `sysfs` file system as follows: • Shared memory partition data: `cat /proc/ppc64/lparcfg` • Virtual I/O bus attributes: `/sys/bus/vio/` directory. • Virtual I/O device attributes: `/sys/bus/vio/devices/` directory. This directory has a subdirectory for each device. Look in the subdirectory for each device to see the virtual I/O device statistics for each device. • Shared Memory statistics: **`amsstat`** (included in powerpc-utils) • Shared Memory graphical monitoring: **`amsvis`** (included in powerpc-utils-python)	• Statistics about the shared memory partition: – I/O entitled memory set for the shared memory partition – Memory weight of the shared memory partition – Amount of physical memory allocated to the shared memory partition – Size of the shared memory pool to which the shared memory partition belongs – Frequency that data is written to the shared memory pool from the paging space device – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the virtual I/O bus, such as the highest amount of physical memory the shared memory partition has ever used for its I/O devices. • Statistics about the virtual I/O devices, such as the frequency that the device tried to map a page to perform an I/O operation and was unable to obtain sufficient memory. In this situation, the attempt fails and delays the I/O operation. • Statistics about the tools: – The packages *powerpc-utils* and *powerpc-utils-python* are user space packages. – The **`amsstat`** script can be run from a Linux logical partition to display shared memory statistics associated with the logical partition. – The **`amsvis`** tool is a python based graphical tool that displays similar information in a graphical manner. This tool is capable of aggregating data from multiple Linux shared memory logical partitions to obtain a picture of cross logical partition performance of shared memory Linux logical partitions.

Related concepts

Factors that influence the performance of shared memory partitions
In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Adjusting the shared memory configuration to improve performance

You can use the Hardware Management Console (HMC) to adjust the configuration of your shared memory environment to improve its performance. For example, you can change the I/O entitled memory or the memory weight that is assigned to a logical partition that uses shared memory (also referred to as a *shared memory partition*).

Determining the I/O entitled memory for a shared memory partition

After you create a new logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) or you dynamically add or remove a virtual adapter, you can use memory statistics that are displayed by the Hardware Management Console (HMC) to dynamically increase and decrease the amount of I/O entitled memory assigned to a shared memory partition.

About this task

The I/O entitled memory set for a shared memory partition needs to be high enough to ensure the progress of I/O operations, and low enough to ensure adequate memory use among all the shared memory partitions in the shared memory pool.

The operating system manages the I/O entitled memory allocated to a shared memory partition by distributing it among the I/O device drivers. The operating system monitors how the device drivers use the I/O entitled memory, and sends usage data to the HMC. You can view the data in the HMC and dynamically adjust the I/O entitled memory that is assigned to a shared memory partition.

For more information about memory settings, see Changing memory settings.

Examples

Creating a new shared memory partition

1. You activate the new shared memory partition. The HMC automatically sets the I/O entitled memory for the shared memory partition.
2. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.
3. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector. (Dynamically decreasing the I/O entitled memory also changes the I/O entitled memory mode to the manual mode.)
4. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

Dynamically adding a virtual adapter to a shared memory partition in the auto I/O entitled memory mode

1. You dynamically add a virtual adapter to a shared memory partition. The HMC automatically increases the I/O entitled memory that is assigned to the shared memory partition.
2. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.
3. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector. (Dynamically decreasing the I/O entitled memory also changes the I/O entitled memory mode to the manual mode.)
4. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

Dynamically adding a virtual adapter to a shared memory partition in the manual I/O entitled memory mode

1. You ensure that the shared memory partition has enough I/O entitled memory to accommodate the new adapter by dynamically increasing the I/O entitled memory of the shared memory partition.
2. You dynamically add a virtual adapter to the shared memory partition.
3. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.

4. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector.
5. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

Dynamically removing a virtual adapter from a shared memory partition

1. You dynamically remove a virtual adapter from a shared memory partition. If the I/O entitled memory mode is in the auto mode, the HMC automatically decreases the I/O entitled memory that is assigned to the shared memory partition.
2. You reset the data collector.
3. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.
4. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector. (If the I/O entitled memory mode is in the auto mode, dynamically decreasing the I/O entitled memory also changes the I/O entitled memory mode to the manual mode.)
5. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

The following example is another way to accomplish this example for AIX shared memory partitions:

1. Determine the amount of physical memory that the virtual adapter (that you plan to remove) currently uses by running the **`lparstat`** command from the AIX command line.
2. If the I/O entitled memory mode is in the auto mode, dynamically change the I/O entitled memory mode to the manual mode by running the **`chhwres`** command from the HMC command line.
3. Using the HMC graphical interface, dynamically remove the virtual adapter.
4. Using the HMC graphical interface, dynamically decrease the I/O entitled memory that is assigned to the shared memory partition by the amount that you identified in step “1” on page 213.

Managing security for logical partitions and operating systems

When all logical partitions are managed by the Hardware Management Console, you can control who has access to the HMC and the system. You can also use the IBM eServer™ Security Planner to help you plan a basic security policy for each of the operating systems on your system.

When all logical partitions are managed by the Hardware Management Console (HMC), the system administrator for the HMC can control who has access to the HMC and the managed systems by creating HMC user roles. The user roles control who can access different parts of the HMC and what tasks they can perform on the managed system.

You can use the IBM eServer Security Planner to help you plan a basic security policy for each of the operating systems on your IBM Power Systems hardware. The planner provides you with a list of recommendations for setting password rules, resource-access rules, logging and auditing rules, and other security settings that are specific to the operating system.

Related information

eServer Security Planner

Troubleshooting IBM i logical partitions

If you have problems with a partitioned system, determine if the problem is specific to logical partitions or is a system problem. If your problem is specific to logical partitions, you can use the reference codes to resolve the error. However, specific recovery actions and tasks might require the assistance of your next level of support.

Debugging network server description error messages for AIX logical partitions

This topic provides a list of network server description (NWSD) error codes and explanations to help you debug NWSD error messages for AIX logical partitions.

You could encounter error messages when you try to vary on an AIX logical partition. These error messages are displayed if you provide information that does not apply to a logical partition running on the server when you create your network server description (NWSD). All error messages related to the NWSD are displayed in QSYSOPR and indicate a description of the problem and a resolution to the problem.

Table 32. NWSD error messages

Reason codes	Code explanations
00000001	*NWSSTG was specified as the IPL source, but no storage space was found.
00000002	The partition specified in the PARTITION parameter was not found.
00000003	The partition specified in the PARTITION parameter is not a GUEST partition (that is, the TYPE parameter for the partition specified in the PARTITION parameter does not have a value of *GUEST).
00000004	There is already an NWSD in the IBM i logical partition that is active and using the partition specified in the PARTITION parameter of the NWSD.
00000005	The partition specified in the PARTITION parameter of the NWSD is powered on (perhaps through the LPAR configuration interface or from another IBM i logical partition.)
00000006	The partition is set to start from a stream file (stmf) and that did not work. You should note that the user performing the vary on operation needs read access to the IPL STMF parameter.
00000007	The NWSD is set to start from a network-storage space (NWSSTG), but the kernel could not found the NWSSTG. Some common reasons are that the storage space does not have a disk partition that is formatted as type 0x41 or is marked as startable.
00000008	The partition would not start. There are a variety of reasons why the partition will not start. You should look at the information for this partition and start reviewing the SRCs.
00000009	The partition identified as the logical partition is not configured. You should specify who has power controlling access to the partition.
00000010	A network server storage space linked to this network server is damaged. Contact your next level of support.
00000011	Contact your next level of support to find a proper solution to the problem.
00000012	The resource name you selected in the RSRCNAME parameter is not valid. Use the Work with Hardware Resources (WRKHDWRSC) command with the TYPE(*CMN) parameter to help determine the resource name.
00000013	The resource you selected in the RSRCNAME command exists, but is not in the partition you specified. Use the WRKHDWRSC command with the TYPE(*CMN) parameter to help determine a resource name in the partition you specified.

Table 32. NWSD error messages (continued)

Reason codes	Code explanations
00000014	Unable to determine partition for resource name. Either specify a partition directly or update the resource definition at the HMC to indicate the client partition.
00000015	Unknown error occurred. Contact your next level of support.

Troubleshooting errors for Linux partitions by using IBM i virtual I/O resources

In many cases, you can troubleshoot and resolve errors specific to Linux logical partitions using IBM i virtual I/O resources without having to call service and support.

Debugging network server description error messages

This topic provides a list of network server description (NWSD) error codes and explanations to help you debug NWSD error messages for Linux logical partitions.

You could encounter error messages when you try to vary on a Linux logical partition. These error messages are displayed if you provide information when you create your network server description (NWSD) that does not apply to a logical partition running on the server. All error messages related to the NWSD are displayed in QSYSOPR indicating a description of the problem and a resolution to the problem.

Table 33. NWSD error messages

Reason codes	Code explanations
00000001	*NWSSTG was specified as the IPL source, but no storage space was found.
00000002	The partition specified in the PARTITION parameter was not found. Use the CHGNWSD IBM i Control Language (CL) command to compare the partition name in the NWSD with the partition name created on the Hardware Management Console (HMC), and change the partition name as necessary.
00000003	The partition specified in the PARTITION parameter is not a GUEST partition (that is, the TYPE parameter for the partition specified in the PARTITION parameter does not have a value of *GUEST).
00000004	There is already an NWSD in the IBM i logical partition that is active and using the partition specified in the PARTITION parameter of the NWSD.
00000005	The partition specified in the PARTITION parameter of the NWSD is powered on (perhaps through the LPAR configuration interface or from another IBM i logical partition.)
00000006	The partition is set to start from a stream file (stmf) and that did not work. You should note that the user performing the vary on operation needs read access to the IPL STMF parameter.
00000007	The NWSD is set to start from a network-storage space (NWSSTG), but the kernel could not found the NWSSTG. Some common reasons are that the storage space does not have a disk partition that is formatted as type 0x41 or is marked as startable.
00000008	The partition would not start. There are a variety of reasons why the partition will not start. You should look at the information for this partition and start reviewing the SRCs.

Table 33. NWSD error messages (continued)

Reason codes	Code explanations
00000009	The partition identified as the logical partition is not configured. You should specify who has power controlling access to the partition.
00000010	A network server storage space linked to this network server is damaged. Contact your next level of support.
00000011	Contact your next level of support to find a proper solution to the problem.
00000012	The resource name you selected in the RSRCNAME parameter is not valid. Use the Work with Hardware Resources (WRKHDWRSC) command with the TYPE(*CMN) parameter to help determine the resource name.
00000013	The resource you selected in the RSRCNAME command exists, but is not in the partition you specified. Use the WRKHDWRSC command with the TYPE(*CMN) parameter to help determine a resource name in the partition you specified.
00000014	Unable to determine partition for resource name. Either specify a partition directly or update the resource definition at the HMC to indicate the client partition.
00000015	Unknown error occurred. Contact your next level of support.

Troubleshooting Linux virtual tape errors

You can troubleshoot and recover from many common errors that are related to Linux virtual tape without having to call service and support.

If errors occur while you access Linux virtual tape, examine the file /proc/systemi/viotape. It describes the mapping between IBM i device names and Linux device names and records the last error for each tape device.

Table 34. Common errors and recovery scenarios for troubleshooting Linux virtual tape

Error	Recovery scenario
Device unavailable	Make sure the device is varied off in the IBM i logical partition.
Not ready	Retry the operation. If the operation still fails with the same description in /proc/iSeries/viotape, verify that the correct medium is in the tape drive.
Load failure or cleaning cartridge found	Verify that the correct medium is in the tape drive.
Data check or Equipment check	Verify that you are using a supported block size to read or write the tape. All known tape devices that are supported by IBM can use a block size of 20 KB (supplied by the -b 40 argument to tar).
Internal error	Contact your service representative.

Situations requiring the assistance of an authorized service provider

Some IBM i troubleshooting tasks on the server require the assistance of an authorized service provider. These tasks are not common and are only performed if the authorized service provider deems it necessary.

If you need to perform any of these tasks on your server, consult the IBM Support Portal website for information on server support.

Main storage dumps on IBM i logical partitions

When your system performs a main storage dump, contact service and support.

On a system with logical partitions, two types of failures can cause main storage dumps: server failure and logical partition failure.

Failures caused by server processing hardware or server firmware might cause the entire server to fail. Software failures in a logical partition cause only that logical partition to fail. A server failure may cause a platform system dump. A logical partition failure may cause a main storage dump only on that logical partition.

You can also force a main storage dump on a logical partition or managed system when you are directed to do so by an authorized service provider.

Using remote service with logical partitions

You can use the Hardware Management Console (HMC) to enable remote services with logical partitions. Remote service is a method that an authorized service provider can use to access your managed system through a modem.

⚠ **Attention:** Use this procedure only when directed to do so by service and support, and ensure that remote service is deactivated when your authorized service provider is finished with it. It is a security risk to leave remote service enabled when not in use. Someone could access your server without your knowledge.

1. Create a user ID.
2. Click **Service Applications** → **Remote Support** → **Customize Inbound Connectivity Settings**.

Shutting down a power domain with logical partitions

You can use the Hardware Management Console (HMC) to power off, repair, and power on the appropriate power domain when a disk unit I/O adapter (IOA) fails. By using this method, you can replace the IOA without restarting the logical partition or managed system.

⚠ **Attention:** Use this procedure only when directed to do so by service and support. Incorrect use of this function can cause loss of data. It can also cause failures that may be incorrectly diagnosed as expensive hardware failures.

When a disk unit IOA fails, communication with the disk units (which is controlled by the IOA) is lost, resulting in a disk unit attention SRC, and possibly partial or complete loss of system responsiveness.

Related information

Performing dumps

Troubleshooting the RMC connection between the logical partition and the HMC

To perform dynamic partitioning operations, you require a Resource Monitoring and Control (RMC) connection between the logical partition and the Hardware Management Console (HMC). If you cannot add or remove processors, memory, or I/O devices to or from a logical partition, check whether the RMC connection is active. Failure of the RMC connection is one of the most common reasons for failure of dynamic partitioning operations.

Before you begin, complete the following procedure:

1. Check the value of the RMC connection state that is cached in the data repository of the HMC by running the following command from the HMC command-line interface:

   ```
   lssyscfg -r lpar -m cec_name -F
   name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
   dlpar_proc_capable,dlpar_io_capable
   ```

The value of the **rmc_state** attribute must either be active or inactive. Also, all the capabilities must be enabled.

For example:

```
#lssyscfg -r lpar -m cec_name -F
name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
dlpar_proc_capable,dlpar_io_capable
lpar01,1,9.5.23.194,1,1,1,1
....
lpar0n,1.9.5.24.###,1,1,1,1
```

If the value of the **rmc_state** attribute or all the capabilities are not set to 1, perform a system rebuild to refresh the data by running the `chsysstate -m system name -o rebuild -r sys` command. If the rebuild operation does not change the value, complete steps 2 and 3.

2. Ensure that the firewall of the HMC is lifted for the RMC port by using the HMC graphical user interface. For the procedure, see solution 1.
3. Ensure that the firewall of the HMC is authenticated for the HMC to receive the request from the logical partition and the logical partition is authenticated to receive the request from the HMC by either using Secure Shell (SSH) or Telnet.

When the operating system on the logical partition is Linux, ensure that Reliable Scalable Cluster Technology (RSCT) Red Hat Package Managers (RPMs) **rsct.core**, **rsct.core.utils**, and **src** are installed. For more information about how to install the RPMs, see Service and productivity tools for SLES on POWER Linux servers for SUSE Linux Enterprise Server operating system and Service and productivity tools For Managed RHEL for Red Hat Enterprise Linux operating system.

The following table lists the steps to check the RMC connection and possible solutions when the connection fails.

Table 35. Steps to check for RMC failure and solutions

Scenario	Solution
Verify whether the firewall settings block the logical partition that is managed by the HMC.	1. To verify the Firewall configuration of the LAN adapter, perform the following steps by using the HMC: a. In the navigation pane, open **HMC Management**. b. In the work pane, click **Change Network Settings**. c. Click the **LAN Adapters** tab. d. Select any LAN adapter other than the `eth0` adapter that connects the HMC with the service processor, and click **Details**. e. On the **LAN Adapter** tab, under **Local area network information**, verify whether **Open** is selected and **Partition communication** status is displayed as enabled. f. Click the **Firewall Settings** tab. g. Ensure that the RMC application is one of the applications that are displayed in **Allowed Hosts**. If it is not displayed in **Allowed Hosts**, select the RMC application under **Available Applications** and click **Allow Incoming**. h. Click **OK**.

Table 35. Steps to check for RMC failure and solutions (continued)

Scenario	Solution
Verify whether the `/tmp` folder in the HMC is 100% full by running the **`df`** command, with superuser privilege.	You must remove unused files in the `/tmp` folder to free up space.

Related information

Checking the status of the management domain and the peer domain

Verifying RMC connections for the mobile partition

RMC network port usage, data flows, and security

Performance considerations for logical partitions

You can manage and enhance the performance of logical partitions so that your system uses its resources in the most efficient manner.

You can manage and enhance the performance of a AIX logical partition by configuring the AIX operating system.

Managing IBM i performance ensures that your managed system is efficiently using resources and that your managed system provides the best possible services to you and to your business. Moreover, effective performance management can help you quickly respond to changes in your managed system and can save on expenses by postponing costly upgrades and service fees.

Related information

AIX Performance management

Performance Tools Guide and Reference

Performance Toolbox Version 2 and 3 Guide and Reference

Power Systems Capacity on Demand

Performance considerations for shared memory partitions

You can learn about performance factors (such as shared memory overcommitment) that influence the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). You can also use shared memory statistics to help you determine how to adjust the configuration of a shared memory partition to improve its performance.

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

A shared memory configuration is considered over committed when the sum of the logical memory that is assigned to all of the shared memory partitions is greater than the amount of physical memory in the shared memory pool.

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

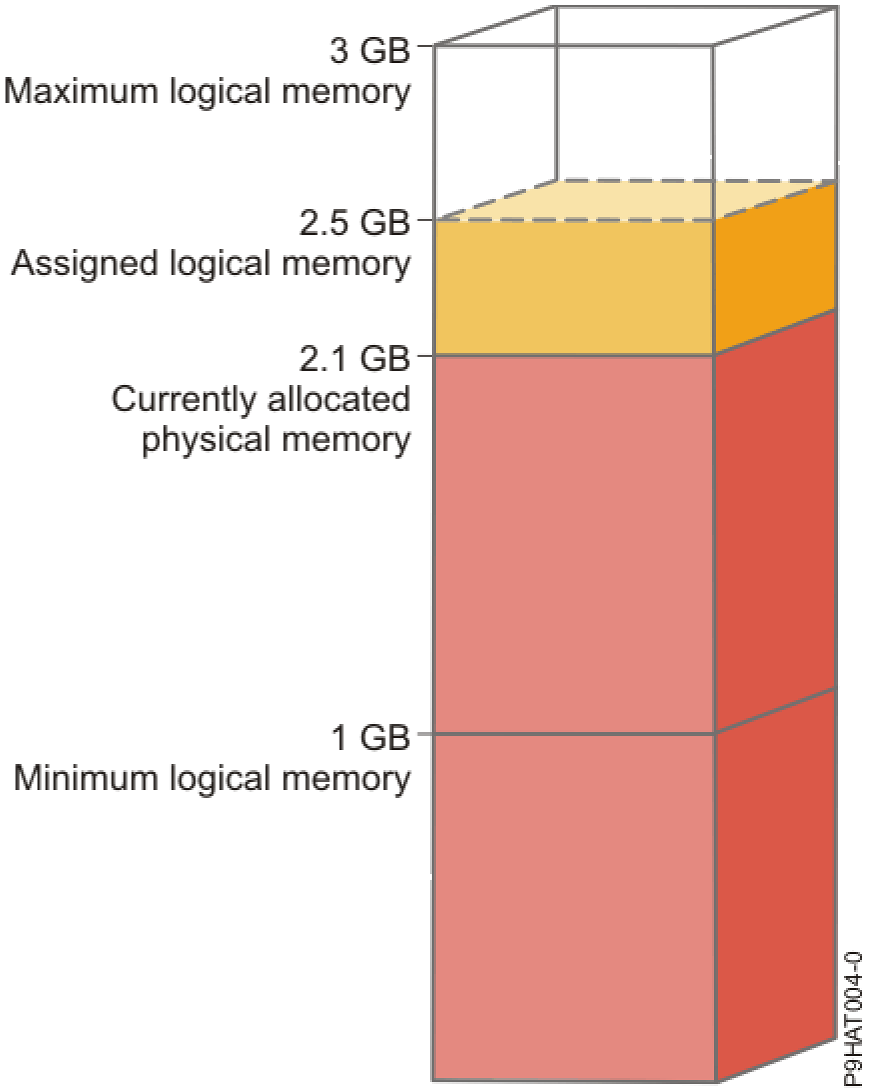

Figure 9. A shared memory partition in a logically overcommitted memory configuration

The figure shows a shared memory partition that is assigned 2.5 GB of logical memory. Its maximum logical memory is 3 GB and its minimum logical memory is 1 GB. The figure also shows that the amount of physical memory that is currently allocated to the shared memory partition from the shared memory pool is 2.1 GB. If the workload that runs in the shared memory partition currently uses 2.1 GB of memory and requires an additional 0.2 GB of memory, and the shared memory pool is logically overcommitted, the hypervisor allocates an additional 0.2 GB of physical memory to the shared memory partition by assigning memory pages that are not currently in use by other shared memory partitions.

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

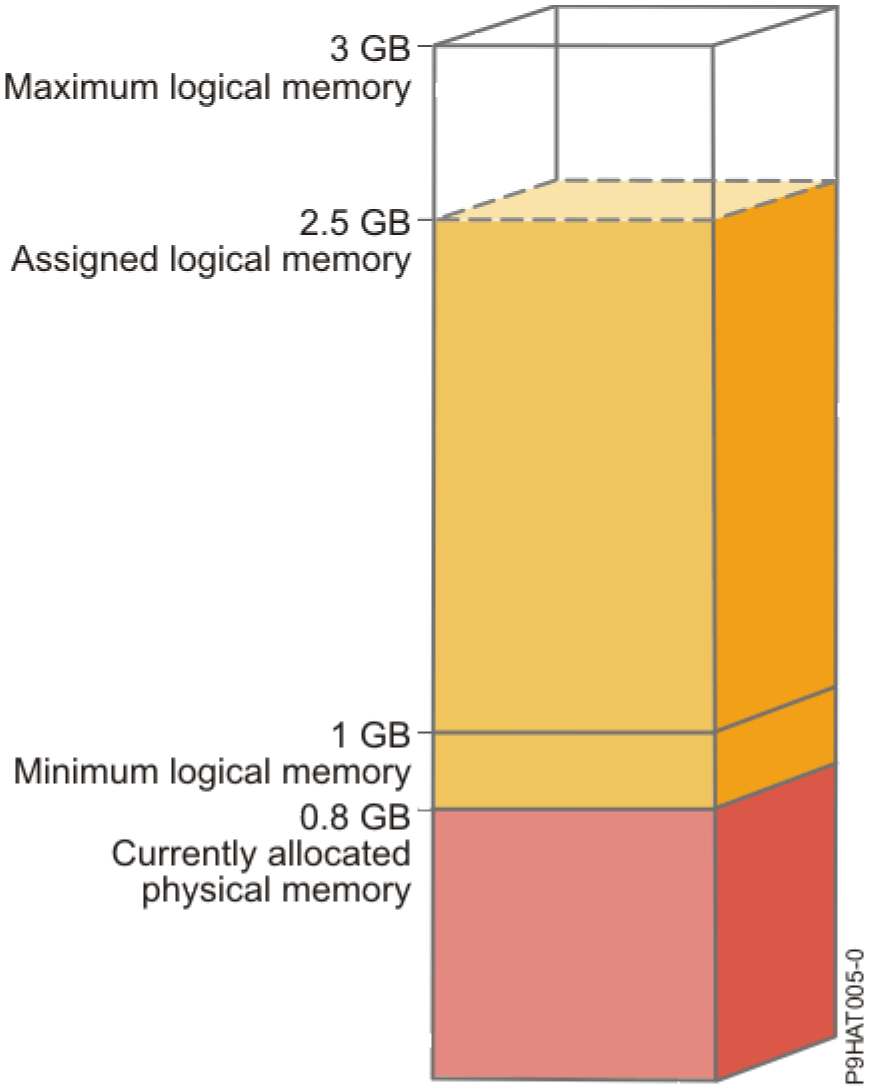

Figure 10. A shared memory partition in a physically over committed memory configuration

The figure shows a shared memory partition that is currently allocated 0.8 GB of physical memory and assigned 2.5 GB of logical memory. If the workload that runs in the shared memory partition currently uses 0.8 GB of memory and requires an additional 1.5 GB of memory, and the shared memory pool is physically over committed, the hypervisor stores 1.5 GB of the shared memory partition's memory in its paging space device.

When the shared memory partition needs to access data on the paging space device, the hypervisor directs a paging VIOS partition to read the data from the paging space device and write the data to the shared memory pool. The more memory that the hypervisor must store on the paging space device, the more often the hypervisor and paging VIOS partition need to read and write data between the paging space device and the shared memory pool. Compared to directly accessing data that is stored in the shared memory pool, it takes more time to access data that is stored in the paging space device. Thus, in general, the less over committed the memory configuration of a shared memory partition, the better its performance.

The operating systems that run in the shared memory partitions help improve the performance of the shared memory partitions with over committed memory configurations by providing the hypervisor with information about how the operating system uses the physical memory that is allocated to it. Using this information, the hypervisor can store data that the operating system accesses the least often in the paging space device and store the data that the operating system accesses the most often in the shared memory pool. This reduces the frequency that the hypervisor needs to access the paging space device and increases the performance of the shared memory partition.

Related concepts

Factors that influence the performance of shared memory partitions
In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Example: A shared memory configuration that is logically overcommitted
When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

Example: A shared memory configuration that is physically overcommitted
When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

Shared memory distribution
The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Related reference

Performance statistics for shared memory
The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

The following table describes the types of workloads that are appropriate to run in shared memory configurations that are logically and physically over committed. It also describes the types of workloads that are not appropriate to run in a shared memory configuration.

Table 36. Workloads to run in logically over committed configurations, physically over committed configurations, and dedicated memory configurations

Workloads for logically overcommitted configurations	Workloads for physically overcommitted configurations	Workloads for dedicated memory configurations
• Workloads that peak at opposite and varying times. • Workloads with memory residency requirements that have a low average. • Workloads that do not have a sustained load. • Logical partitions that serve as failover and backup logical partitions when configured on the same server as their primary counterparts. • Test and development environments.	• Workloads that run the AIX operating system and use the file cache. • Print servers, file servers, network applications, and other workloads that are less sensitive to I/O latency. • Workloads that are inactive most of the time.	• Workloads with high quality of service criteria. • Workloads that consistently use memory resources due to sustained peak load. • High-performance computing (HPC) workloads.

In addition to the degree to which the memory configuration of a shared memory partition is over committed, the following factors can influence the performance of a shared memory partition:

- The workload that runs in a shared memory partition, the number of virtual adapters that are assigned to the shared memory partition, and the I/O entitled memory set for the shared memory partition all directly affect the performance of I/O devices. These factors can cause I/O devices to operate at their minimum memory requirements rather than their optimal memory requirements. This can cause delays in I/O operations.
- The amount of I/O entitled memory that is required for optimal performance depends on the workload and number of adapters configured.
- The operating systems that run in shared memory partitions cannot use memory affinity. Some applications rely on memory affinity to improve their performance.
- The shared memory partition might be suspended if it attempts to access data on its paging space device when the following situations occur simultaneously:
  - The paging VIOS partition becomes unavailable. For example, you shut down the paging VIOS partition or the paging VIOS partition fails.
  - The shared memory partition is not configured to use redundant paging VIOS partitions to access its paging space device.

Related concepts

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Related reference

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Where to view statistics	Statistics to view
HMC utilization data	• Statistics about the shared memory pool, such as: – Size of the shared memory pool – Total amount of memory that is over committed – Total amount of logical memory that is assigned to the shared memory partitions – Total amount of I/O entitled memory that is assigned to the shared memory partitions – Total amount of physical memory that the shared memory partitions currently use for their I/O devices – Amount of memory from the shared memory pool that the hypervisor uses to manage the shared memory partitions – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the shared memory partitions, such as: – Amount of logical memory assigned to the shared memory partition – Amount of physical memory from the shared memory pool that is allocated to the shared memory partition – Amount of memory that is over committed – I/O entitled memory assigned to the shared memory partition – Amount of physical memory that the shared memory partition currently uses for its I/O devices – Memory weight of the shared memory partition

Where to view statistics	Statistics to view
IBM i See IBM® i to view shared memory statistics in IBM i.	• Statistics about the shared memory pool, such as: – Total number of page faults for all of the shared memory partitions – Total time, in milliseconds, that the processors waited for page faults to be resolved – Total physical memory, in bytes, that is assigned to the shared memory pool – Sum of the logical memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the I/O entitled memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the physical memory, in bytes, that the shared memory partitions that are active currently use for their I/O devices • Statistics about the shared memory partition, such as: – Memory weight of the shared memory partition – Amount of physical memory, in bytes, from the shared memory pool that is currently used by the shared memory partition – Number of times that the shared memory partition waited for a page fault – The time, in milliseconds, that the shared memory partition waited for page faults to be resolved – Maximum amount of memory, in bytes, that the shared memory partition can assign to data areas that are shared between the operating system and the server firmware – I/O entitled memory assigned to the shared memory partition – Minimum amount of physical memory, in bytes, required for all of the configured I/O devices to operate – Optimal amount of physical memory, in bytes, required for I/O devices to maximize throughput performance – Amount of physical memory, in bytes, that the shared memory partition currently uses for its I/O devices – Highest amount of physical memory, in bytes, that the shared memory partition has used for its I/O devices since the last time the shared memory partition was activated or since the last time the memory statistics were reset, whichever is most recent – Number of delayed I/O operations since the last time the shared memory partition was activated

Where to view statistics	Statistics to view
Linux View memory statistics for Linux in the `sysfs` file system as follows: • Shared memory partition data: `cat /proc/ppc64/lparcfg` • Virtual I/O bus attributes: `/sys/bus/vio/` directory. • Virtual I/O device attributes: `/sys/bus/vio/devices/` directory. This directory has a subdirectory for each device. Look in the subdirectory for each device to see the virtual I/O device statistics for each device. • Shared Memory statistics: **`amsstat`** (included in powerpc-utils) • Shared Memory graphical monitoring: **`amsvis`** (included in powerpc-utils-python)	• Statistics about the shared memory partition: – I/O entitled memory set for the shared memory partition – Memory weight of the shared memory partition – Amount of physical memory allocated to the shared memory partition – Size of the shared memory pool to which the shared memory partition belongs – Frequency that data is written to the shared memory pool from the paging space device – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the virtual I/O bus, such as the highest amount of physical memory the shared memory partition has ever used for its I/O devices. • Statistics about the virtual I/O devices, such as the frequency that the device tried to map a page to perform an I/O operation and was unable to obtain sufficient memory. In this situation, the attempt fails and delays the I/O operation. • Statistics about the tools: – The packages *powerpc-utils* and *powerpc-utils-python* are user space packages. – The **`amsstat`** script can be run from a Linux logical partition to display shared memory statistics associated with the logical partition. – The **`amsvis`** tool is a python based graphical tool that displays similar information in a graphical manner. This tool is capable of aggregating data from multiple Linux shared memory logical partitions to obtain a picture of cross logical partition performance of shared memory Linux logical partitions.

Related concepts

Factors that influence the performance of shared memory partitions
In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Performance considerations for over committed shared memory partitions
Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Adjusting the shared memory configuration to improve performance

You can use the Hardware Management Console (HMC) to adjust the configuration of your shared memory environment to improve its performance. For example, you can change the I/O entitled memory or the memory weight that is assigned to a logical partition that uses shared memory (also referred to as a *shared memory partition*).

Troubleshooting the RMC connection between the logical partition and the HMC

To perform dynamic partitioning operations, you require a Resource Monitoring and Control (RMC) connection between the logical partition and the Hardware Management Console (HMC). If you cannot add or remove processors, memory, or I/O devices to or from a logical partition, check whether the RMC connection is active. Failure of the RMC connection is one of the most common reasons for failure of dynamic partitioning operations.

Before you begin, complete the following procedure:

1. Check the value of the RMC connection state that is cached in the data repository of the HMC by running the following command from the HMC command-line interface:

   ```
   lssyscfg -r lpar -m cec_name -F
   name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
   dlpar_proc_capable,dlpar_io_capable
   ```

   The value of the **`rmc_state`** attribute must either be active or inactive. Also, all the capabilities must be enabled.

   For example:

   ```
   #lssyscfg -r lpar -m cec_name -F
   name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
   dlpar_proc_capable,dlpar_io_capable
   lpar01,1,9.5.23.194,1,1,1,1
   ....
   lpar0n,1.9.5.24.###,1,1,1,1
   ```

   If the value of the **`rmc_state`** attribute or all the capabilities are not set to 1, perform a system rebuild to refresh the data by running the `chsysstate -m system name -o rebuild -r sys` command. If the rebuild operation does not change the value, complete steps 2 and 3.
2. Ensure that the firewall of the HMC is lifted for the RMC port by using the HMC graphical user interface. For the procedure, see solution 1.
3. Ensure that the firewall of the HMC is authenticated for the HMC to receive the request from the logical partition and the logical partition is authenticated to receive the request from the HMC by either using Secure Shell (SSH) or Telnet.

When the operating system on the logical partition is Linux, ensure that Reliable Scalable Cluster Technology (RSCT) Red Hat Package Managers (RPMs) **`rsct.core`**, **`rsct.core.utils`**, and **`src`** are installed. For more information about how to install the RPMs, see Service and productivity tools for SLES on POWER Linux servers for SUSE Linux Enterprise Server operating system and Service and productivity tools For Managed RHEL for Red Hat Enterprise Linux operating system.

The following table lists the steps to check the RMC connection and possible solutions when the connection fails.

Table 37. Steps to check for RMC failure and solutions

Scenario	Solution
Verify whether the firewall settings block the logical partition that is managed by the HMC.	1. To verify the Firewall configuration of the LAN adapter, perform the following steps by using the HMC: a. In the navigation pane, open **HMC Management**. b. In the work pane, click **Change Network Settings**. c. Click the **LAN Adapters** tab. d. Select any LAN adapter other than the `eth0` adapter that connects the HMC with the service processor, and click **Details**. e. On the **LAN Adapter** tab, under **Local area network information**, verify whether **Open** is selected and **Partition communication** status is displayed as enabled. f. Click the **Firewall Settings** tab. g. Ensure that the RMC application is one of the applications that are displayed in **Allowed Hosts**. If it is not displayed in **Allowed Hosts**, select the RMC application under **Available Applications** and click **Allow Incoming**. h. Click **OK**.
Verify whether the `/tmp` folder in the HMC is 100% full by running the **`df`** command, with superuser privilege.	You must remove unused files in the `/tmp` folder to free up space.

Related information

Checking the status of the management domain and the peer domain

Verifying RMC connections for the mobile partition

RMC network port usage, data flows, and security

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled “Cookies, Web Beacons and Other Technologies”.

Programming interface information

Logical partitioning publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.2, IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.4, and IBM Virtual I/O Server Version 3.1.2.

Trademarks

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Setting up the virtualization environment

Note

Before using this information and the product it supports, read the information in "Notices" on page 39.

This edition applies to IBM® AIX® Version 7.2, to IBM AIX Version 7.1, to IBM AIX Version 6.1, to IBM i 7.4 (product number 5770-SS1), to IBM Virtual I/O Server Version 3.1.2 and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

Contents

Setting up the virtualization environment

A template is a collection of configuration preferences that can be reused and quickly applied to multiple targets. You can use templates to set up your virtualization environment. Templates simplify the deployment process because templates contain many of the settings that you previously configured by using the Hardware Management Console (HMC) command-line interface or the HMC graphical user interface (GUI) version 8.1.0, or earlier.

The templates functions are supported only when a server is managed by the HMC, or when a server is co-managed by the HMC and PowerVM® NovaLink, with the HMC in the master mode.

The PowerVM NovaLink architecture enables management of highly scalable cloud deployment by using the PowerVM technology and OpenStack solutions. The architecture provides a direct OpenStack connection to a PowerVM server. The NovaLink partition runs the Linux® operating system and the partition runs on a server that is virtualized by PowerVM. The server is managed by PowerVC or other OpenStack solutions.

There are two types of templates: system template and partition template. You can use system templates to define system configuration settings that include general system properties and virtual environment settings. You can use partition templates to specify logical partition settings that include general partition properties, processor and memory configuration, virtual networks and virtual storage configuration, logical Host Ethernet Adapters and logical single root I/O virtualization (SR-IOV) port settings. Templates do not contain target-specific information. Therefore, you can use templates to configure any system or partition in your environment.

Templates are further classified into quick-start templates or user-defined templates.

Quick-start templates are contained in the `template` folder that is accessible by using the template library. You cannot edit the quick-start templates, however you can copy and change them to suit your requirements.

User-defined templates are the templates that you create. User-defined templates contain configuration details that are specific to your environment. You can create a user-defined template by using any of the following methods:

- Copy an existing template and modify the new template according to the requirements of your environment.
- Capture the configuration details of a currently running system or partition and save the details into a new template.

What's new in Setting up the virtualization environment

Read about new or changed information in Setting up the virtualization environment since the previous update of this topic collection.

November 2020

- The following topics were updated with information about the partition keystore feature and enhancements to the HMC graphical user interface:
  - "Changing a system template" on page 5
  - "Changing a partition template" on page 24
  - "Creating a logical partition by using a template" on page 33
  - "Creating logical partitions by using Create partition option" on page 36

May 2020

- The following topics were updated with information about enhancements to single root I/O virtualization (SR-IOV) logical ports:
  - “Changing a partition template” on page 24
  - “Creating a logical partition by using a template” on page 33

October 2019

- The following topics were updated with information about enhancements to single root I/O virtualization (SR-IOV) logical ports:
  - “Changing a partition template” on page 24
  - “Creating a logical partition by using a template” on page 33
- The following topics were updated with information about persistent memory support:
  - “Changing a partition template” on page 24
  - “Creating a logical partition by using a template” on page 33

May 2019

- The following topics were updated with information about enhancements to the HMC graphical user interface:
  - “Importing a system template” on page 21
  - “Importing a partition template” on page 31
  - “Creating a logical partition by using a template” on page 33
  - “Creating logical partitions by using Create partition option” on page 36
- The following topics were updated with information about RDMA over Converged Ethernet (RoCE) support:
  - “Changing a system template” on page 5
  - “Deploying a system by using a system template” on page 16
  - “Changing a partition template” on page 24
  - “Creating a logical partition by using a template” on page 33

August 2018

- The following topics were updated for the secure boot feature:
  - “Changing a system template” on page 5
  - “Changing Virtual I/O Server settings” on page 10
  - “Changing a partition template” on page 24
  - “Creating a logical partition by using a template” on page 33
- The following topics were updated for changes to the template capture and deploy operations:
  - “Capturing a system configuration” on page 4
  - “Deploying a system by using a system template” on page 16
  - “Capturing a partition configuration” on page 23
  - “Creating a logical partition by using a template” on page 33
- The following topic was updated for support of USB devices for VIOS installation:
  - “Deploying a system by using a system template” on page 16

Accessing the template library

All templates reside in the template library, which is accessible from the Hardware Management Console (HMC).

About this task

To view and select the templates that are available in the template library, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Select the **System** tab to view the existing system templates, or the **Partition** tab to view the existing partition templates.
4. Select a template from the list of templates that are displayed.

Results

You can view, modify, deploy, copy, import, export, or delete user-defined templates that are available in the template library. You cannot edit the quick-start templates, however you can edit a copy of the quick-start template.

System templates

System templates contain configuration information about resources such as system properties, shared processor pools, reserved storage pool, shared memory pool, physical I/O adapters, Host Ethernet Adapters, single root I/O virtualization (SR-IOV) adapters, Virtual I/O Server, virtual networks, virtual storage, and initial program load (IPL).

The single root I/O virtualization (SR-IOV) specification defines extensions to the PCI Express (PCIe) specification. SR-IOV allows virtualization of the physical ports of an adapter so that the ports can be shared by multiple partitions that are running simultaneously. For example, a single physical Ethernet port appears as several separate physical devices. To share the ports of an SR-IOV capable adapter, the adapter must first be enabled for the SR-IOV shared mode. After an adapter is enabled for SR-IOV shared mode, SR-IOV logical ports can be assigned to logical partitions.

Many of the system settings that you previously configured by using the Hardware Management Console (HMC) command-line interface or the HMC graphical user interface (GUI) version 8.1.0, or earlier can now be completed by using the **Deploy System from Template** wizard. For example, you can configure the Virtual I/O Servers, virtual network bridges, and virtual storage settings when you use the wizard to deploy a system by using a system template.

The template library includes quick-start system templates, which contain configuration settings based on common usage scenarios. Quick-start system templates are available for your immediate use.

You can also create user-defined system templates that contain configuration settings that are specific to your environment. You can create a user-defined template by copying any template that is available in the template library and then changing it to suit your requirements. You can also capture the configuration of an existing system and save the details in a template. You can deploy that template to other systems that require the same configuration.

System templates are primarily used to deploy settings to new systems. To deploy new systems, complete the following tasks:

1. “Viewing system template configuration information” on page 4
2. “Prerequisites for deploying a system by using a system template” on page 16
3. “Capturing a system configuration” on page 4 (optional)

4. “Deploying a system by using a system template” on page 16

You can also complete the following tasks by using system templates:

- “Changing a system template” on page 5
- “Copying a system template” on page 20
- “Importing a system template” on page 21
- “Exporting a system template” on page 21
- “Deleting a system template” on page 22

Viewing system template configuration information

Before you deploy a system template on a system, you must review the configuration details of the template to determine whether you want to use a quick-start template or create a user-defined template. Until you create one or more user-defined templates, quick-start templates are the only templates that are available in the template library.

About this task

To view the configuration information by using the Hardware Management Console (HMC), complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. In the Templates and O/S Images window, click the **System** tab.
4. Select the system template that you want to view and click **Actions** > **View**.

   You can view the details of **Physical I/O**, **Host Ethernet Adapter**, **SR-IOV**, **Virtual I/O Servers**, **Virtual Networks**, **Virtual Storage**, **Shared Processor Pool**, **Shared Memory Pool and Reserved Storage**, and **Advanced System Settings** by clicking the relevant tabs that are displayed. Alternately, you can view the template details from the **Deploy System Template** wizard.
5. Click **Close**.

Capturing a system configuration

Capturing a system gathers the current configuration of the system that is in the running state and includes information about the Virtual I/O Server (VIOS), virtual network, virtual storage, and system settings. You can capture these details of a running system and save the information as a user-defined system template by using the Hardware Management Console (HMC). This function is useful if you want to deploy multiple systems with the same configuration. If you want to use a quick-start template, you need not complete this task.

About this task

To capture the configuration of a running system by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
   a) Click **All Systems**. The All Systems page is displayed.
   b) In the work pane, select the system and click **Actions** > **View System Properties**.

      The Properties page is displayed. You can choose only one system at a time.

c) Expand **System Actions** > **Templates** > **Capture Configuration as Template** > **with Physical I/O**, for capturing the configuration with physical I/O information. The option to capture physical I/O information is available only when the system is in the running state. Alternatively, for capturing the configuration without physical I/O information, expand **System Actions** > **Templates** > **Capture Configuration as Template** > **without Physical I/O**. Information about the system configuration such as Virtual I/O Servers, Virtual Networks, and Virtual Storage is displayed in the **Template Details** page. All non-target specific data is contained in the appropriate fields of the system template.

2. In the **Capture as System Template** page, specify the name for the template file in the **Template Name** field.
3. Enter a description for the template in the **Template Description** field and click **OK** to save the captured template, or click **Cancel** if you want to cancel the operation.

   In the **Capture as System Template** page, you can also view the progress of the capture operation. A message indicates the successful completion of the capture operation. Appropriate warning or error messages are displayed when applicable and errors result in a failure of the capture operation.

Results

The template is available in the template library. You can deploy a system by using this template, or you can modify any aspect of the template before using the template to deploy a system.

Related tasks

Deploying a system by using a system template
You can deploy systems by using system templates that are available in the template library of the Hardware Management Console (HMC). The **Deploy System from Template** wizard guides you to provide specific information about the target system that is required to complete the deployment on the selected system.

Changing a system template
You can change the details that are specified in a captured, or user-defined system template and save the changes in a new system template. You can also overwrite the template by saving the changes in the same template. You can use this template to deploy other systems by using the Hardware Management Console (HMC).

Changing a system template

You can change the details that are specified in a captured, or user-defined system template and save the changes in a new system template. You can also overwrite the template by saving the changes in the same template. You can use this template to deploy other systems by using the Hardware Management Console (HMC).

About this task

To change the system template by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. To change the Physical I/O settings, click the **Physical I/O** tab. You can enable or disable that uses the captured I/O information. Click the **Use Captured I/O Information** check box to use the captured physical I/O information.

When you enable or disable the **Use Captured I/O Information** check box, the check box is displayed in the read-only mode in the **Hardware Virtualized I/O** tab for the system configuration, and in the **Hardware Virtualized I/O** tab of the **Add VIOS** wizard.

6. To change the Host Ethernet Adapter (HEA) or single root I/O virtualization (SR-IOV) settings, click the **Hardware Virtualized I/O** tab.

   a) In the **Host Ethernet Adapter** tab, you can change the **HEA Port Group Settings** and the **HEA Physical Port Settings**. You can add an HEA by clicking the **Add** tab, and to remove an HEA, select the HEA and click the **Remove** tab.

      The HEA tab contains general adapter-level settings that are applied to all discovered physical HEA ports and HEA port groups during deployment of the system.

   b) In the **SR-IOV** tab, the table of SR-IOV adapters displays the properties of the hardware virtualized I/O adapters that are available.

      If the captured system template contains information about RDMA over Converged Ethernet (RoCE) logical ports, that information about the RoCE logical ports is listed in the table of SR-IOV adapters. If you are not using any captured I/O information, you must specify the adapter settings during deployment of the system. You can view the physical Ethernet port settings for the selected adapters. Select a physical location code from the **Physical Port** area to view the physical Ethernet port settings. You can view the **Speed**, **Flow Control**, and the **MTU Size** of the Ethernet adapter. You can also change the values of the **Label** and **Sub-label** fields.

7. To change the Virtual I/O Server (VIOS) settings, click the **Virtual I/O Servers** tab. Select the VIOS that you want to rename. You can specify the name in the **VIOS Name** field. To add a VIOS, complete the following steps:

   a) Click the **Add VIOS tab**.

   b) In the **General** tab, you can specify the name for the VIOS in the **VIOS Name** field.

   c) Select a value for the **Boot Mode** field.

   d) To enable synchronization of the current profile, select the **Save Configuration Change to profile** check box.

      When this option is selected, the partition profile is always synchronized with the last activated partition profile.

   e) In the **Advanced Settings** area of the **General** tab, you can select or clear the **Automatic Start With Managed System**, **Mover Service Partition**, **Enable Connection Monitoring**, **Enable Redundant Error Path Reporting**, **Enable Time Reference**, **Enable VTPM**, and **Allow Performance Information Collection** fields. You can select a value for the **Secure Boot** field if you are using HMC Version 9.1.920, or later, and when the firmware is at level FW920, or later. If the HMC is at a Version 9.2.950, or later, and when the firmware is at level FW950, or later, you can specify a value of either 0 kilobytes (KB) or a value within the range of values supported by the system for the **Keystore Size** field.

   f) In the **Processor** tab, if you select **Shared** for the processor mode, you can set the processor weight as capped or uncapped. When you set the processor weight as uncapped, you must specify a value for the processor weight in the **Weight** field.

   g) In the **Virtual Processors** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   h) In the **Advanced Settings** area, you can select a value for the **Processor Compatibility Mode**.

   i) In the **Processor** tab, if you select **Dedicated** as the processor mode, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields in the **Processors** area.

   j) In the **Advanced Settings** area, you can select a value for the **Processor Compatibility Mode** and the **Idle Processor Sharing** fields.

   k) In the **Memory** tab, if you select **Shared** for the memory mode, you can set the value for the memory in either MB or GB.

   l) In the **Memory Allocation** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

m) In the **Memory** tab, if you select **Dedicated** for the memory mode, you can set the value for the memory in either MB or GB.

n) In the **Memory Allocation** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

The **Advanced Settings** area is displayed only when you use dedicated memory. You can enable **Enable Memory Expansion**, and the **Huge Page Memory**.

- If you enable **Enable Memory Expansion**, you can specify a value in the range 1.0 - 10.0 for the Active Memory Expansion (AME) factor.
- If you enable **Huge Page Memory**, you can specify values for **Minimum**, **Allocated**, and **Maximum** fields.

o) In the **Physical I/O Adapters** tab, if you select the **Use Captured I/O Information** check box, you can view the details of the physical I/O adapters that were captured.

p) Click the **Hardware Virtualized I/O** tab and then click the **SR-IOV** tab. If you select the **Use Captured I/O Information** check box, you can view the details of the captured physical Ethernet port. If the captured system template contains information about RDMA over Converged Ethernet (RoCE) logical ports, that information about the RoCE logical ports is also listed. If you do not use the captured I/O information, you can select the logical port and the Shared Ethernet Adapter (SEA) backing device that you want to assign to the VIOS. Also, details about the RoCE logical ports are displayed if they are available in the system, but you cannot use RoCE logical ports as SEA backing devices.

q) If you do not select a SEA backing device and click the **Advanced Settings** tab, you can select available values for the **OS MAC Address Restrictions**, **VLAN ID Restrictions**, **Port VLAN ID**, and the **802.1Q Priority** fields.

r) To remove a logical port, select the logical port to be removed and click **Remove Selected**.

s) Click the **HEA** tab. You can select the logical port and the Shared Ethernet Adapter (SEA) backing device that you want to assign to the VIOS.

t) To add an HEA, click the **Add** tab. In the **HEA Port Group Settings** area, you can select a value for the multi-core Scaling (MCS) Value field. In the **HEA Physical Port Settingss** area, you can set the speed, select full duplex or half duplex mode, enable or disable flow control, and specify the maximum receiving packet size.

u) To remove an HEA, select the HEA to be removed and click **Remove Selected**.

v) To delete a VIOS, select the VIOS, then right-click to select **Remove VIOS**.

8. To change the virtual network settings, click the **Virtual Networks** tab. You can change the details of a specific virtual network, virtual switch, or network bridge by selecting the table row of the virtual network, virtual switch, or network bridge, then right-click on the selected entry. You can also delete a network, switch, or network bridge. To add a virtual network, complete the following steps:

a) Click **Add Virtual Network**.

b) Click the **Network Name** tab.

c) In the **Virtual Network Settings** area, enter a value for the virtual network in the **Virtual Network Name** field.

d) Select a value for the **Virtual Network Type**.

If you selected **Internal Network** for the **Virtual Network Type** field, enter a value in the **Virtual Network ID** field. Optionally, if you selected **Bridged Network**, you must select a value for the **IEEE 802.1q Tagging**.

e) Select the **Add new virtual network to all Virtual I/O Servers** check box to assign the virtual network to all the Virtual I/O Servers that are specified in the template.

f) In the **Virtual Switch Settings** area, click **Use an existing virtual switch** to use existing virtual switches, or click **Create a new virtual switch**.

If you clicked **Create a new virtual switch**, enter a value for the name of the virtual switch in the **Virtual Switch Name** field.

g) Click **Next**.

h) If the template had any existing virtual network bridges, you can click **Select existing Virtual Network Bridge**. Optionally, you can also click **Create a New Virtual Network Bridge**.

i) In the **Virtual Network Bridge Settings** area, enter a value for the **Virtual Network Bridge PVID (PowerVM)** field.

j) Select a value for the **Failover** field. If you choose to use failover for VIOS, you must also select a value for the **Secondary VIOS** and the **Load Sharing** fields.

k) Select a value for the **Primary VIOS** field.

l) In the **Optional Settings** area, select values for the **Jumbo Frame**, **Large Send**, and **QoS** fields.

m) Click **Next** to view and edit values in the **Load Sharing** tab only if you are using an existing virtual network bridge.

Clicking **Next** will allow you Otherwise, clicking **Next** will display the **Summary** tab.

n) In the **Load Sharing** tab, click **Use an existing Load Sharing Group** to use an existing load sharing group, or click **Create a new Load Sharing Group**.

o) If you select **Create a new Load Sharing Group**, enter a value for the **New Load Group PVID** field.

p) Click **Next**.

q) In the **summary** tab, a summary of the configuration that was selected for the virtual network is displayed. You can review the configuration details and click **Finish** to add the virtual network to the VIOS specified in the template.

9. To change the virtual storage settings, click the **Virtual Storage** tab. You can change the details of the **Shared Storage Pool Clusters**. You can assign each VIOS that is listed in the template to an actual shared storage pool cluster that is managed by the HMC. You can specify a **Media Repository** for each VIOS.
10. To change the shared processor pool settings, click the **Shared Processor Pool** tab. You can add shared processor pools, rename the pools (except the Default Pool), and adjust the processing units that are assigned to each pool. You can also delete a shared processor pool.
11. To change the Shared Memory Pool and Reserved Storage Pool settings, click the **Shared Memory Pool and Reserved Storage Pool** tab. You can change details of a Shared Memory Pool such as the size and maximum pool size. You can also specify whether Active Memory Deduplication must be enabled. You can also change the settings of the Reserved Storage Device Pool. You can select a single VIOS, or specify redundancy settings by selecting **Redundant VIOS**.
12. To change the advanced system settings, click the **Advanced System Settings** tab. You can change details of the **Power On/Off Configuration**, and the **Memory and performance Configuration** fields. Click **Save and Exit** to save the changes in the same template, or select **Save As** to save the changes in a new template. Click **Cancel** to exit without making any changes.

Changing physical I/O settings

You can change the system I/O settings of a system by changing a quick-start or captured system template by using the Hardware Management Console (HMC). This changed template is used to deploy the system.

About this task

To change the physical I/O settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.

3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. To change the Physical I/O settings, click the **Physical I/O** tab. You can enable or disable using captured I/O information. Click the **Use Captured I/O Information** check box to use the captured physical I/O information.
6. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Changing Hardware Virtualized I/O settings

You can change the Hardware Virtualized I/O settings that are specified in a system template and either overwrite the template or save the changes in a new template by using the Hardware Management Console (HMC). You can change the Host Ethernet Adapter (HEA) and single root I/O virtualization (SR-IOV) settings.

About this task

To change HEA and SR-IOV settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. Click the **Hardware Virtualized I/O** tab.
6. In the **Host Ethernet Adapter** tab, you can modify the **HEA Port Group Settings** and the **HEA Physical Port Settings**.
   a) In the **HEA Port Group Settings** area, you can select a value for the **Multi-Core Scaling (MCS) Value**. You can complete this task for each port group that is listed.
   b) In the **HEA Physical Port Settings** area, for each of the ports that are listed, you can set the speed, specify full duplex or half duplex mode, enable or disable flow control, and specify the maximum receiving packet size.
   c) To add an HEA, click the **Add** tab.
   d) In the **HEA Port Group Settings** area, you can select a value for the **Multi-Core Scaling (MCS) Value**.
   e) In the **HEA Physical Port Settings** area, you can set the speed, specify full duplex or half duplex mode, enable or disable flow control, and specify the maximum receiving packet size.
   f) To remove an HEA, select the HEA to be removed and click the **Remove** tab.
7. In the **SR-IOV** tab, you must specify the SR-IOV physical Ethernet ports settings when you deploy the template.
8. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Results

The HEA tab contains general adapter-level settings that are applied to all discovered physical HEA ports and HEA port groups during deployment. For example, if you set the Port Group MCS Value to 2, the value for Maximum Logical Ports changes. During deployment, each port group that is located on any Host Ethernet Adapter on the target system has a Port Group MCS Value of 2 and makes the appropriate number of logical ports available. Similarly, the four physical port settings apply to each physical port on all the HEAs discovered during deployment.

Changing Virtual I/O Server settings

You can change the Virtual I/O Server (VIOS) settings that are specified in a system template by using the Hardware Management Console (HMC). You can add or remove VIOS, change the properties of the VIOS, or change the resources that are assigned to the VIOS.

About this task

To change the VIOS settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. To change the Virtual I/O Server (VIOS) settings, click the **Virtual I/O Servers** tab. Select the VIOS that you want to rename. You can specify the name in the **VIOS Name** field. To add a VIOS, complete the following steps:

   a) Click the **Add VIOS tab**.

   b) In the **General** tab, you can specify the name for the VIOS in the **VIOS Name** field.

   c) Select a value for the **Boot Mode** field.

   d) To enable synchronization of the current profile, select the **Save Configuration Change to profile** check box.

      When this option is selected, the partition profile is always synchronized with the last activated partition profile.

   e) In the **Advanced Settings** area of the **General** tab, you can select or clear the **Automatic Start With Managed System**, **Mover Service Partition**, **Enable Connection Monitoring**, **Enable Redundant Error Path Reporting**, **Enable Time Reference**, **Enable VTPM**, and **Allow Performance Information Collection** fields. You can select a value for the **Secure Boot** field if you are using HMC Version 9.1.920, or later, and when the firmware is at level FW920, or later. If the HMC is at a Version 9.2.950, or later, and when the firmware is at level FW950, or later, you can specify a value of either 0 kilobytes (KB) or a value within the range of values supported by the system for the **Keystore Size** field.

   f) In the **Processor** tab, if you select **Shared** for the processor mode, you can set the processor weight as capped or uncapped. When you set the processor weight as uncapped, you must specify a value for the processor weight in the **Weight** field.

   g) In the **Virtual Processors** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   h) In the **Advanced Settings** area, you can select a value for the **Processor Compatibility Mode**.

   i) In the **Processor** tab, if you select **Dedicated** as the processor mode, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields in the **Processors** area.

   j) In the **Advanced Settings** area, you can select a value for the **Processor Compatibility Mode** and the **Idle Processor Sharing** fields.

   k) In the **Memory** tab, if you select **Shared** for the memory mode, you can set the value for the memory in either MB or GB.

   l) In the **Memory Allocation** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   m) In the **Memory** tab, if you select **Dedicated** for the memory mode, you can set the value for the memory in either MB or GB.

   n) In the **Memory Allocation** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

      The **Advanced Settings** area is displayed only when you use dedicated memory. You can enable **Enable Memory Expansion**, and the **Huge Page Memory**.

      - If you enable **Enable Memory Expansion**, you can specify a value in the range 1.0 - 10.0 for the Active Memory Expansion (AME) factor.
      - If you enable **Huge Page Memory**, you can specify values for **Minimum**, **Allocated**, and **Maximum** fields.

   o) In the **Physical I/O Adapters** tab, if you select the **Use Captured I/O Information** check box, you can view the details of the physical I/O adapters that were captured.

   p) Click the **Hardware Virtualized I/O** tab and then click the **SR-IOV** tab. If you select the **Use Captured I/O Information** check box, you can view the details of the captured physical Ethernet port. If the captured system template contains information about RDMA over Converged Ethernet (RoCE) logical ports, that information about the RoCE logical ports is also listed. If you do not use the captured I/O information, you can select the logical port and the Shared Ethernet Adapter (SEA) backing device that you want to assign to the VIOS. Also, details about the RoCE logical ports are displayed if they are available in the system, but you cannot use RoCE logical ports as SEA backing devices.

   q) If you do not select a SEA backing device and click the **Advanced Settings** tab, you can select available values for the **OS MAC Address Restrictions**, **VLAN ID Restrictions**, **Port VLAN ID**, and the **802.1Q Priority** fields.

   r) To remove a logical port, select the logical port to be removed and click **Remove Selected**.

   s) Click the **HEA** tab. You can select the logical port and the Shared Ethernet Adapter (SEA) backing device that you want to assign to the VIOS.

   t) To add an HEA, click the **Add** tab. In the **HEA Port Group Settings** area, you can select a value for the multi-core Scaling (MCS) Value field. In the **HEA Physical Port Settingss** area, you can set the speed, select full duplex or half duplex mode, enable or disable flow control, and specify the maximum receiving packet size.

   u) To remove an HEA, select the HEA to be removed and click **Remove Selected**.

   v) To delete a VIOS, select the VIOS, then right-click to select **Remove VIOS**.

6. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Changing virtual network settings

You can change the virtual network settings that are specified in a system template and overwrite the template or save the changes in a new template by using the Hardware Management Console (HMC). You can change the virtual switches or the properties of the network bridge. You can also add or remove virtual networks.

About this task

To change the virtual network settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon.

2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. Click the **Virtual Networks** tab.
6. In the **Virtual Networks** area, right-click the virtual network that you want to change.

   You can change virtual network name, and the **Load Balance Group**. You can also delete the virtual network.
7. In the **Virtual Switches** area, right-click the virtual network that you want to change.

   You can change virtual switch name and specify whether the switching mode is Virtual Ethernet Bridging (VEB) or Virtual Ethernet Port Aggregator (VEPA).
8. In the **Virtual Bridges** area, right-click the network bridge that you want to change.

   You can change the network bridge name, specify whether failover and load balancing is enabled, and specify the VIOS that are associated with the network bridge.
9. In the **Advanced Settings** area, you can specify the Quality of Service (QoS) Priority, and enable **Jumbo Frame** and **Large Send**.
10. Click **OK**.
11. To add a virtual network, complete the following steps:

    a) Click **Add Virtual Network**.

    b) Click the **Network Name** tab.

    c) In the **Virtual Network Settings** area, enter a value for the virtual network in the **Virtual Network Name** field.

    d) Select a value for the **Virtual Network Type**.

    If you selected **Internal Network** for the **Virtual Network Type** field, enter a value in the **Virtual Network ID** field. Optionally, if you selected **Bridged Network**, you must select a value for the **IEEE 802.1q Tagging**.

    e) Select the **Add new virtual network to all Virtual I/O Servers** check box to assign the virtual network to all the Virtual I/O Servers that are specified in the template.

    f) In the **Virtual Switch Settings** area, click **Use an existing virtual switch** to use existing virtual switches, or click **Create a new virtual switch**.

    If you clicked **Create a new virtual switch**, enter a value for the name of the virtual switch in the **Virtual Switch Name** field.

    g) Click **Next**.

    h) If the template had any existing virtual network bridges, you can click **Select existing Virtual Network Bridge**. Optionally, you can also click **Create a New Virtual Network Bridge**.

    i) In the **Virtual Network Bridge Settings** area, enter a value for the **Virtual Network Bridge PVID (PowerVM)** field.

    j) Select a value for the **Failover** field. If you choose to use failover for VIOS, you must also select a value for the **Secondary VIOS** and the **Load Sharing** fields.

    k) Select a value for the **Primary VIOS** field.

    l) In the **Optional Settings** area, select values for the **Jumbo Frame**, **Large Send**, and **QoS** fields.

    m) Click **Next** to view and edit values in the **Load Sharing** tab only if you are using an existing virtual network bridge.

    Clicking **Next** will allow you Otherwise, clicking **Next** will display the **Summary** tab.

    n) In the **Load Sharing** tab, click **Use an existing Load Sharing Group** to use an existing load sharing group, or click **Create a new Load Sharing Group**.

    o) If you select **Create a new Load Sharing Group**, enter a value for the **New Load Group PVID** field.

p) Click **Next**.

q) In the **summary** tab, a summary of the configuration that was selected for the virtual network is displayed. You can review the configuration details and click **Finish** to add the virtual network to the VIOS specified in the template.

12. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Changing virtual storage settings

You can change the virtual storage settings that are specified in a system template and either overwrite the template or save the changes in a new template by using the Hardware Management Console (HMC). You can change the Virtual I/O Server (VIOS) that belong to a shared storage pool, or you can add or remove media repositories.

About this task

To change the virtual storage settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. Click the **Virtual Storage** tab.
6. In the **Virtual Shared Storage Pool Clusters** area, select the shared storage pool cluster that must be assigned to the VIOS. Alternatively, you can select **Choose at Deploy**. Complete this task for all the Virtual I/O Servers that are listed.
7. In the **Media Repositiries** area, you can configure the media repository for each of the VIOS listed in the **Virtual Shared Storage Pool Clusters** area. In the **Media Repository Size** field, specify the size in percentage values. The unit of measurement can be GB or MB.
8. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Changing shared processor pool settings

You can change the shared processor settings that are specified in a system template and either overwrite the template or save the changes in a new template by using the Hardware Management Console (HMC). This template can be used to deploy a system with the changed shared processor pool settings. You can configure a maximum of 63 shared processor pools.

About this task

To change the shared processor pool settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

The **Template Detail** page is displayed.

5. Click the **Shared Processor Pool** tab.
6. In the **Pool Name** field, you can specify a name for the pool. You can specify the reserved and maximum processing units that must be assigned to each processor pool in the **Reserved Processing Units** and **Maximum Processing Units** fields.

   This adds another row to the table.

   Note: You cannot rename the default pool.
7. Click **Add Another** tab to add another shared processor pool and specify the name of the pool you want to add. This adds another row to the table.
8. To remove a shared processor pool, select the pool that you want to remove and click **Remove**.
9. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Changing Shared Memory Pool and Reserved Storage Device Pool settings

You can change the Shared Memory Pool and Reserved Storage Device Pool settings that are specified in a system template and either overwrite the template or save the changes in a new template by using the Hardware Management Console (HMC). This template can be used to deploy a system with the changed shared memory pool and reserved storage settings.

About this task

To change the Shared Memory Pool and Reserved Storage Device Pool settings of a template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. Click the **Shared Memory Pool and Reserved Storage** tab.
6. In the **Shared Memory Pool** area, you can specify a value for the pool size in the **Pool Size** field in GB or MB.
7. In the **Maximum Pool Size** field, you can specify a value for the maximum pool size in GB or MB.
8. From the **Active memory Deduplication** list, you can enable or disable Active memory Deduplication.
9. In the **Reserved Storage Pool** area, you can specify the names of the Virtual I/O Server (VIOS) to be used as the **First Paging VIOS** and the **Second Paging VIOS**.
10. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Changing advanced system settings

You can change the advanced system settings that are specified in a captured or user-defined system template and overwrite the template, or save the changes in a new template by using the Hardware Management Console (HMC). You can change the system performance, and Power On or Power Off policies for the systems. This reduces the time that is taken to deploy multiple systems.

Changing system performance settings

Your can change your system performance settings for the system to efficiently use its resources. Effective performance management can help you quickly respond to changes in your system and save on expenses by postponing upgrades and service fees. You can also improve performance by setting appropriate values for the Logical Memory Block (LMB), Huge Page Count, and also by specifying the Power On and Power Off configuration.

About this task

To change the system performance settings, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. Click the **Advanced System Settings** tab.

   a) In the **Memory and Performance Configuration** area, you can specify the **Logical Memory Block (LMB)** size.

   b) In the **Huge Page Count** field, specify a value.
6. Click **Save and Exit** to overwrite the changes made in the template, or click **Save As** to save the changes in a new template.

Changing power policy settings

You can change the power policy settings that are specified in a system template and overwrite the template or save the changes in a new template by using the Hardware Management Console (HMC). You can use this template to deploy a system with the changed power policy settings.

About this task

To change the system power policy settings, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and select the system template that you want to change.
4. Click **Action** > **Edit**.

   The **Template Detail** page is displayed.
5. Click the **Advanced System Settings** tab.

   a) In the **Power On/Off Configuration** area, select a value from the **Server Firmware Start Policy** list.

   b) From the **System Power Off Policy** list, select a value.

   c) From the **Power On Speed** list, select a value.

   d) You can enable or disable **Auto Power Restart**.
6. Click **Save and Exit** to overwrite the changes that are made in the template, or click **Save As** to save the changes in a new template.

Prerequisites for deploying a system by using a system template

Review the prerequisites before you deploy a system by using templates.

The **Deploy System from Template** wizard guides you through the deployment operation by using a system template. The wizard includes the following tasks:

- Selecting a system template when you start the deployment from a system, or selecting a system when you start the deployment from a template library.
- Configuring the system settings, assigning I/O adapters, and creating Virtual I/O Servers.
- Installing the Virtual I/O Server (VIOS) software.
- Configuring the network and storage I/O settings.

A system can be in one of the following states before you deploy a system template. Before you deploy a system template, review the following information to understand the impact of continuing with the deployment of the system when the system is in one of the following states:

- The system is in the manufacturing default configuration. You can start the system deployment on the system.
- The system is not in the manufacturing default configuration, and there are no partitions. If you attempt to deploy a system template on an already configured system, the HMC displays a warning message. If you click **OK**, the deployment continues and any previous partition configuration is removed. The system is configured with the partitions that are specified in the system template.
- The system is not in the manufacturing default configuration and has partitions. If you attempt to deploy a system template, the HMC displays a warning message. If you start the system deployment, existing logical partition configuration data is removed.

You must back up your data to perform a system recovery, when needed.

Related information

Backing up and recovering data

Deploying a system by using a system template

You can deploy systems by using system templates that are available in the template library of the Hardware Management Console (HMC). The **Deploy System from Template** wizard guides you to provide specific information about the target system that is required to complete the deployment on the selected system.

Before you begin

Before you deploy a system, consider the following prerequisites:

- The HMC is at Version 8.1.0, Service Pack 1, or later.
- The hypervisor is in the operating or standby state.
- The managed system is in the operating or standby state.
- The managed system does not have any logical partition that is associated to it.

  Note: If logical partitions are already configured on the managed system, a warning message is displayed. If you continue with the deployment, the HMC completes the following actions:

  - All system level configurations will be initialized or set to default values.
  - All the logical partitions that are in the running state will be shutdown and removed automatically.
  - All the Virtual I/O Servers that are in the running state will be shutdown and removed automatically.
- If you install the VIOS from a Network Installation Management (NIM) server, you must have the NIM server information that is required by the HMC.

About this task

When you deploy a system from a template, the HMC checks whether the configuration specified in the chosen template suits the required system capabilities.

To deploy a system by using the system template, complete the following steps:

Note: During the deployment, you can view all of the configuration settings specified in the template by clicking the **Template Details** tab.

Procedure

1. In the navigation pane, click the **Resources** icon .
   a) Click **All Systems**. The All Systems page is displayed.
   b) In the work pane, select the system and click **Actions** > **View System Properties**. The Properties page is displayed. You can choose only one system at a time.
   c) Expand **System Actions** > **Templates** > **Deploy System from Template**. When you deploy a template on the system, a system data check operation is performed on the system that you selected. To check the status of a system, select a server from the list of servers and click **Reset**.
2. If the check operation displays warning or error messages, you can select another template from the template library. If the chosen template is compatible with the target system and no warning or error messages are displayed, or when the displayed warning message is acceptable, click **Next** to continue with the deployment. If the target system has logical partitions, you will be prompted with a message that indicates that the logical partition and the current system settings will be deleted which cannot be recovered and whether you want to continue with the deployment. Click **Yes** to continue with the deployment, or **No** to exit the wizard.

   Alternatively, you can also select a system and click **Templates** > **Template Library**. Go to step 5.
3. If you choose to deploy the template form the template library, click the **System** tab and select a template from the list of templates. Click **Reset**. If the check operation displays warning or error messages, you can select another template from the template library. If the chosen template is compatible with the target system and no warning or error messages are displayed, or when the displayed warning message is acceptable, click **Next** to continue with the deployment. If the target system has any logical partitions, you will be prompted with a message that indicates that the logical partitions will be deleted and whether you want to continue with the deployment. Click **Yes** to continue with the deployment, or **No** to exit the wizard.
4. In the **SR-IOV Adapter Settings** page, you can select the single root I/O virtualization (SR-IOV) capable adapter.

   You can assign a logical port from an SR-IOV adapter that is in shared mode to a Virtual I/O Server (VIOS). By default, the adapters are in the dedicated mode.
   a) Click **Shared** to change the mode to the shared mode.
   b) Click **Configure**.
   c) In the **SR-IOV Adapter Settings** area, you can view or change the settings of each physical port on the SR-IOV adapter.
   d) Click **Next**.
5. If the chosen template is compatible with the target system, the **VIOS Configuration Summary** page is displayed. You can optionally change the VIOS name. Click **Next**.
6. In the **Physical I/O** page, complete the following steps:
   a) In the **Physical I/O Adapters** area, you can choose one or more Virtual I/O Servers to which you want to assign physical I/O adapters. You can view the adapters that are available in other drawers of the system by selecting the drawer in the **View adapters in** field. You can choose not to assign any physical I/O adapters to the VIOS by selecting the **Unassigned** option.

      Note:

- The HMC communicates with the target system and provides a list of physical I/O adapters that can be assigned to a VIOS. You can assign each adapter to a single VIOS, you need not assign all of the adapters.
- If you are using a captured system template from HMC V9.1.910, or earlier, and if the template contains captured physical I/O adapter information, the physical I/O adapter information captured from the system during the capture operation is not used by the HMC. The HMC checks for the available physical I/O adapters on the destination servers that can be assigned to the Virtual I/O Servers. The available physical I/O adapters are listed in the **Physical I/O** page, and you can choose a physical I/O adapter for assignment to the Virtual I/O Servers from this list. This behavior is also applicable when you do not want to use the captured I/O information (not selecting the **Use Captured I/O Information** check box in the **Physical I/O** page).
- If the template that you use for deployment does not contain any VIOS, the wizard displays only the **System Configuration Progress**, and the **Configuration Summary** pages. The **Configuration Summary** page displays read-only information about the settings as defined in the template. You can review these settings, and click **Next** to view the **System Configuration Progress** page to start the deployment process.
- If PowerVM is already configured on the system (the factory default configuration with all resources assigned to the logical partition, a message is displayed that indicates that a partition is already present in the system. You must reset the system manually before restarting the system deployment.
- Slots that contain cable cards are non-partitionable and cannot be assigned to the VIOS. Therefore, the Physical I/O adapter configuration page does not display slots that contain cable cards although the cable cards are associated with the system.

b) In the **Hardware Virtualized I/O** area, you can assign logical ports to a physical port of an SR-IOV capable adapter. Details about the adapters that are available such as the adapter type are displayed. RDMA over Converged Ethernet (RoCE) logical ports are displayed if RoCE logical ports are available.

- If you use captured I/O information, this page displays the list of adapters assigned for the configuration.
- If you are not using captured I/O information, you can select a value for the **Physical Port, Label, Sub-label** field and you can specify a value for the **Capacity** field. You can complete this step for each of the Virtual I/O Servers that are listed.

c) Click **OK**.

The sum of the percentage capacity values for all the configured logical ports on a physical port must be less than or equal to 100%. To minimize the configuration effort when you add more logical ports, you can reserve some capacity for additional logical ports.

7. In the **System Configuration Progress** page, when you click **Start**, the system configuration starts and you can view the progress of the system configuration and a message that indicates successful configuration is displayed upon completion.
8. When the system update completes and when the VIOS partition is created, you can click the **Next** tab to install the VIOS image.

   Note: This step might take some time, especially when you have to restart the system.
9. In the **VIOS Installation Configuration** page, you can select the **Installation Method**. You can also change the configuration values of the VIOS partitions that are listed. By clicking **Advanced Settings**, you can change the adapter speed, adapter duplex, VLAN tag priority, and the VLAN tag identifier default settings. Click **Next**.

   Note: You can install the VIOS from a NIM server, a management console image, a manual console session, or from an image on a USB device. The fields that require data vary depending on the installation method that you choose. The following options are available depending on the installation method:

   - When you install the VIOS from a NIM server, you must specify the server IP address. The HMC must be able to connect to the NIM server.

- When you install the VIOS from an image repository, you must specify the HMC IP address and the VIOS image name.
- When you install the VIOS from a management console, you must specify the boot mode.
- When you install the VIOS from a USB device, ensure that the USB device has the VIOS image that you want to install. You must specify the HMC IP address, and select the VIOS image from the list of VIOS images. All VIOS images that are save in the USB device are listed.

After you select the installation method, you must also specify an install adapter, the port number, and the VIOS IP address, subnet mask, and default gateway. Additionally, you can view the MAC address of the system when you select the NIM server installation option. You can optionally change the adapter speed, adapter duplex, VLAN tag priority, and the default settings of the VLAN tag identifier by clicking **Advanced Settings**. You can complete this step for each of the listed Virtual I/O Servers.

10. In the **VIOS Installation Progress** page, when you click **Start**, the VIOS software is installed on the system. To view the progress of the VIOS installation, click **Monitor vterm**. A message that indicates successful configuration is displayed when the installation is complete.
11. After the VIOS image is installed and after the RMC connection is established for all the Virtual I/O Servers, you can review the license agreement and click **Accept all VIOS Licenses** to accept the VIOS license agreement.
12. In the **VIOS Network Bridge Configuration** page, you can change the values of the listed network bridges. Click **Next**.

    Note: A network bridge represents the Shared Ethernet Adapter (SEA) and trunk adapter that services a set of externally visible virtual networks. For redundant networks, the network bridge represents the matched set of Shared Ethernet Adapters and trunk adapters on both of the Virtual I/O Servers.

    In the **NetBridge** area of the **Network Configuration** page, you can view a table that contains the available network adapters and ports that are assigned to the installed Virtual I/O Servers. Each VIOS that is installed and is associated with a Network Bridge in the template, has a separate table. You can select at least one port to create the Shared Ethernet Adapter for that VIOS, or select more than one physical port per VIOS, or choose to create a Link Aggregation Device from the ports that are selected on a VIOS. A link aggregation device also known as an EtherChannel device, is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters that are aggregated then act as a single Ethernet device. Link aggregation provide more throughput over a single IP address than a single Ethernet adapter. When you are using a captured template for the system deployment, the ports and **Create Link Aggregation Device** might already be selected.
13. In the **VIOS Virtual Storage Configuration** page, you can associate a VIOS to a shared storage pool. You can configure the Reserved Storage Device Pool, and the media repository volume group. Click **Next**.

    Note: You can assign a VIOS to a shared storage pool cluster or assign it later. A shared storage pool cluster provides distributed storage access to the VIOS partitions in the cluster. You can also configure a Reserved Storage Device Pool. A Reserved Storage Device Pool has reserved storage devices that are also called paging space devices and is similar to a Shared Memory Pool of memory size 0 bytes.

    When you configure a reserved storage pool, the **Deploy System from Template** wizard page displays the available Reserved Storage Devices. You can select an available device from the list of devices to create the Reserved Storage Device Pool. You must select which VIOS should be the first and second paging VIOS. A paging VIOS is a VIOS partition that is assigned to the Shared Memory Pool and provides access to the paging space devices for the logical partitions that are assigned to the Shared Memory Pool.

    In the Media Repository Volume Groups area, an editable field that contains the Media Repository name and a table that contains the available storage devices to assign to the Volume Groups is displayed. You can also configure a Media Repository.

14. In the **I/O Progress page** page, when you click **Start**, the configuration process starts and you can view the I/O configuration. You can click **Next**, after you see a message that indicates successful installation.
15. In the **Summary** page, you can view a summary of the changes. Click **Finish**.

    Your system is now fully deployed based on the configuration settings that were specified in the template.

    Note: If the configuration is unsuccessful, you must exit the **Deploy System from Template** wizard and restart the system deployment. You can exit the wizard by clicking **Finish**.

    - You cannot deploy an incomplete template.
    - If the deploy operation fails immediately after a machine data reset, all the current configuration of the target system is destroyed and you cannot restore the system to the previous state.
    - If the deploy operation fails, the deploy system template wizard creates a VIOS and a message is displayed that indicates that the deployment competed with errors. The VIOS that was created cannot be rolled back. You must clean up the deployment manually or use the **Manage PowerVM** functionality available in the HMC to assign the network or storage to the VIOS that was created.

Recovering from a system deployment failure

If a system deployment by using the system template fails, use the Hardware Management Console (HMC) to reset the system to a nonpartitioned configuration. The factory reset mode (or manufacturing default configuration) is equivalent to the initial single partition configuration of the managed system as received from your service provider. After you reset the system, run the **Deploy System Template** wizard again.

If the system deployment by using the system template fails, the system is not restored or backed up to its previous state. You must manually configure the system by using the HMC command-line interface, or start a new deployment. If the deployment of a system fails, exit the Deploy System Template wizard. Reset the system to a nonpartitioned configuration and restart the deployment process. To reset the system, type the **`rstprofdata`** command from the HMC command line. Specify a value of 4 for the *restore type* parameter. Then, restart the Deploy System Template wizard. The **`rstprofdata`** command removes only the data disk, but the boot disk is retained.

When the system deployment fails during the I/O adapter configuration, network configuration, or virtual storage configuration, you can exit the wizard and complete the configuration by using the **Manage PowerVM** functionality available in the HMC.

Related information

Resetting the managed system to a nonpartitioned configuration

rstprofdata

Copying a system template

You can copy a quick-start or captured system template into a new system template along with the configuration details that are specified in the template by using the Hardware Management Console (HMC).

About this task

To copy a system template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.

3. Click the **System** tab and then select the system template that you want to copy and click **Action** > **Copy**.
4. In the **Copy System Template** page, specify the name for the template in the **Template name** field.

   If a template with the same name exists, the copy fails and an error message is displayed.
5. Click **OK**.

Importing a system template

You can import a system template into the template library by using the Hardware Management Console (HMC).

About this task

Before importing a system template, consider the following restrictions:

- If the system template schema is different from the schema that is supported by the HMC, for example, if a tag that is not a part of the template OpenDocument Spreadsheets (ODS) file element is used, the system template cannot be imported. However, if you are using HMC V9.1.930, or later, you can select another system template and import that system template.
- If the system template file size exceeds 10 MB, the system template cannot be imported and the operation fails. However, if you are using HMC V9.1.930, or later, you can select another system template and import that system template.

To import a system template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab and click **Import**.

   The following restrictions apply when you import a system template:
4. In the **Import System Template** page, click **Browse** to navigate to the required template file.

   After you select the file, the selected file name is displayed in the **Template name** field. You can optionally change the name of the file. If a template with the same name exists, the import operation fails and an error message is displayed. If you are using HMC V9.1.930, or later, you can change the name of the file or select another system template and import that system template.
5. Click **OK**.

Exporting a system template

You can export a system template from the template library by using the Hardware Management Console (HMC).

About this task

To export a system template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab. Select the template and click **Action** > **Export**.

A browser-generated window opens where you can choose to save the exported file.

4. Click the **Save file** tab and specify the file name to which the exported file must be saved.
5. Click **OK**.

Deleting a system template

You can delete a system template from the template library by using the Hardware Management Console (HMC).

About this task

To delete a system template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **System** tab. Select the template and click **Action** > **Delete**.
4. In the **Delete Template** page, click **Yes** to delete the selected template, or click **No** to close the **Delete Template** page.

Partition templates

Partition templates contain details about partition resources, such as physical adapters, virtual networks, and storage configuration. You can create client partitions from the quick-start templates that are available in the template library or from your own user-defined templates. You can use the **Deploy Partition template** wizard to create AIX, IBM i, or Linux logical partitions.

In previous releases, partitions were associated with profiles, which stored the configuration information for that partition. It was only possible to turn on a partition after you activated the partition by choosing a profile.

With the Hardware Management Console (HMC) Version 8.1.0, Service Pack 1, or later, when you create a partition by using a template, a default profile for that partition is created automatically. The profile is based on the configuration that is specified in the template that was used to create the partition. After you create a partition by using a template, the template does not retain any association with the partition that you created. You need not use a template to create a new partition, however, by using templates you can simplify the partition creation process. Templates offer more flexibility than profiles because you can choose from the following options when you create a partition by using a template:

- **Create partition** - Creates a partition based on the template you chose, but does not turn on the partition.
- **Create and activate partition** - Creates a partition based on the template you chose and commits the resources that are associated with that template to the partition. Unlike the **Create partition** option, this option turns on the partition.

The quick-start partition templates included in the template library contain configurations based on common scenarios. However, you can also create user-defined templates that contain configuration settings that are specific to your environment.

Partition templates are primarily useful for creating new partitions. The process of deploying a partition by using a template includes the following tasks:

1. “Prerequisites for creating a logical partition by using a template” on page 23
2. “Viewing partition template details” on page 23 (optional)
3. “Capturing a partition configuration” on page 23 (optional)
4. “Creating a logical partition by using a template” on page 33

You can also complete the following tasks by using partition templates:

- “Changing a partition template” on page 24
- “Copying a partition template” on page 31
- “Importing a partition template” on page 31

Prerequisites for creating a logical partition by using a template

Review the prerequisites before you create a logical partition by using a template.

You can create an AIX, IBM i, or Linux logical partition by using any of the partition templates from the template library. The **Create Partition from Template** wizard guides you through the procedure of creating a logical partition.

The system must be in the running state before you create a logical partition from a template on that system. You cannot create a partition from a template when the system is in the powered-off state.

You can choose only one template or system at a time. The system that you chose to deploy to, or the template name that you chose from the template library is displayed on the screen.

Viewing partition template details

Before you create a logical partition by using a template, review the details in that template. By reviewing the configuration details, you can determine whether that template suits the requirements of your environment. If you are already certain which template you want to use to create a logical partition, this task is optional.

About this task

To view the details of a partition template by using the Hardware Management Console (HMC), complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab and select the partition template that you want to view.
4. Click **Action** > **View**.

   You can view the details such as the processor, memory, physical I/O adapter and the general properties of the partition. You can view the **Virtualization Capabilities** area of the general properties tab to verify whether the partition supports the simplified remote restart feature. You can also view the **Virtual Networks**, **Virtual NICs**, **Virtual Storage**, and **Hardware Virtualized I/O** details of the partition by clicking the relevant tabs that are displayed. Alternately, you can view the template details from the **Create Partition from Template** wizard.

Capturing a partition configuration

You can capture the configuration details from a running partition or from a partition that is not activated and save the configuration as a custom template. You can use this function to create multiple partitions with the same configuration. If you want to use a quick-start template, you need not complete this task.

About this task

To capture the current configuration of a running logical partition by using the Hardware Management Console (HMC), complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
   a) Click **All Systems**. The All Systems page is displayed.
   b) In the work pane, select the system on which the partition is located and click **Actions** > **View System Partitions**. All the partitions that are available on the system are displayed.
   c) Select the partition for which you want to capture the configuration information and click **Actions** > **Templates** > **Capture Partition as a Template**. Details about the partition configuration such as processors, memory, physical I/O adapters, and virtualized I/O adapters are displayed in the Template Details page. If the configuration is captured for a partition that supports the simplified remote restart capability, the captured template displays the feature as enabled in the **Virtualization Capabilities** area of the general properties tab. All data that is not specific to a target is captured in the appropriate fields of the partition template.
2. In the **Capture as Partition Template** page, specify the name for the template file in the **Template Name** field.
3. Enter a description for the template in the **Template Description** field and click **OK** to save the captured template, or click **Cancel** if you want to cancel the operation.

   In the **Capture as Partition Template** page, you can also view the progress of the capture operation. A message indicates the successful completion of the capture operation. Appropriate warning or error messages are displayed when applicable and errors result in a failure of the capture operation.

Results

If you chose to save the template, your custom template is now available in the template library. You can create a partition by using this template. For instructions, see "Creating a logical partition by using a template" on page 33. You can also change the configuration details of the template. For instructions, see "Changing a partition template" on page 24.

Changing a partition template

You can change the details that are specified in a user-defined or captured partition template and save the changes in a new partition template. You can also overwrite the template by saving the changes in the same template. You can use the template to create a logical partition by using the Hardware Management Console (HMC).

About this task

To change the partition template by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab and select the partition template that you want to change.
4. Click **Action** > **Edit**.
5. To change the partition properties of the template, click the **Properties** tab.

   In the **Overview** area of the **General** tab, you can change the general properties such as the partition name and the type of partition. Also, you can view the virtual serial number that is captured in the template. You can enable or disable Simplified Remote Restart, in the **Virtualization Capabilities** area of the **General** tab. You can choose one of the following three options for the value of the **Simplified Remote Restart** field:

- When you select the value **Enabled**, the HMC validates whether the server supports the simplified remote restart capability. If the server supports the feature, the partition is created successfully during partition creation by using the **Create Partition from Template** wizard. When the server does not support the feature, creation of the partition fails and an error message is displayed.
- When you select the value **Enable if Possible**, the HMC validates whether the server supports the simplified remote restart capability. If the server supports the feature, the partition is created successfully during partition creation by using the **Create Partition from Template** wizard. Otherwise, the **Create Partition from Template** wizard completes successfully without the simplified remote restart capability.
- When you select the value **Disabled**, the partition is created during partition creation by using the **Create Partition from Template** wizard without the simplified remote restart capability.

In the **Advanced settings** area, you can configure, enable and disable advanced AIX, Linux, or IBM i features. You can also disable the Live Partition Mobility feature for an AIX, Linux, or IBM i partition. The **Advanced** settings that are displayed depend on the type of partition you selected. You can select a value for the **Secure Boot** field if you are using HMC V9.1.920, or later and when the firmware is at level FW920, or later. Additionally, if the HMC is at a Version 9.2.950, or later, and when the firmware is at level FW950, or later, you can specify a value of either 0 kilobytes (KB) or a value that is supported by the managed system for the **Keystore Size** field. Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

6. To change the shared processor settings of the template, click the **Processor** tab and select Shared for the **Processor Mode** field.

   a) From the **Shared Processor Pool** field, select the shared processor pool for the partition.

   b) Select **Capped** or **Uncapped** as the processor weight. For capped processor weight, specify a value for the weight in the **Weight** field.

   c) In the **Virtual Processors** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   d) In the **Processing Units** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   e) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

7. To change the dedicated processor settings of the template, click the **Processor** tab and select Dedicated for the **Processor Mode** field.

   If you select **Dedicated** as the processor mode, the options for specifying the virtual processors, and processing units, and selecting capped or uncapped processor weight are not available.

   a) In the **Processors** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   b) Click the **Advanced Settings** tab to change the **Processor Compatability Mode** or enable or disable the **Dedicated Donor Mode**.

   c) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

8. To change the shared memory settings of the template, click the **Memory** tab and select Shared for the **Memory Mode** field.

   a) You can select MB or GB as the memory unit.

   b) In the **Memory Allocation** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   c) Click the **Advanced Settings** tab to change the advanced memory settings for the logical partition. From the **Assigned I/O Entitled memory** list, select **Auto** or **Manual**. If the operating system environment is IBM i, you can use **Huge Page Memory**. If the operating system environment is AIX, you can also choose to use active memory expansion by selecting **Enable Active memory Expansion**.

d) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

9. To change the dedicated memory settings of the template, click the **Memory** tab and select Dedicated for the **Memory Mode** field.

   a) You can select MB or GB as the memory unit.

   b) In the **Memory Allocation** area, you can specify values for the **Maximum**, **Allocated**, and **Minimum** fields.

   c) Click the **Advanced Settings** tab to change the advanced memory settings for the logical partition. If the operating system environment is IBM i, you can choose use **Huge Page Memory**. If the operating system environment is AIX, you can also choose to use active memory expansion by selecting **Enable Active memory Expansion**.

   If the processor mode is dedicated, then you can set the memory only to the dedicated mode.

   d) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

10. To change the physical I/O settings of the template, click the **Physical I/O Adapters** tab.

    When you use a template with captured I/O information, the **Physical I/O Adapters** tab displays a table of the captured I/O adapter information and the descriptions of those adapters as captured from the original system. These descriptions might not display the actual adapter type to which the location codes are mapped on the target system. You cannot change the I/O adapter settings when the captured I/O information matches that of the target system. If you do not want to use the captured information, clear the **Use Captured I/O Information** check box. You can assign I/O adapters that are displayed to the partition.

11. To change the configuration of the persistent memory volume in the template, click the **Persistent Memory** tab.

    Persistent memory is supported only when you are using HMC V9.1.940, or later and when the firmware is at level FW940, or later.

    a) In the **Persistent Memory** area, you can view all the virtual persistent memory volumes that are configured for the logical partition.

    If you are using the captured I/O information, you can view the name of persistent memory volume, affinity, and the size of the persistent memory volume that are captured from the original system. If you are not using the captured I/O information, you can complete the following tasks:

    - To add persistent memory volumes, click **Add** in the **Persistent Memory** page.
    - To remove persistent memory volumes, select the persistent memory volumes in the **Persistent Memory** page and click **Action** > **Remove**.
    - You can select the **Affinity** check box in the **Partition Properties** page when you want the operating system to get information about the memory allocation across multiple dual in-line memory modules (DIMMs). This information might be useful for applications that run on the logical partition.

12. To change the virtual network settings of the template, click the **Virtual Networks** tab in the **Details** area.

    a) In the **Partition Virtual Networks** area, you can choose either **Choose Virtual Networks during Deployment** or **Specify Virtual Networks in this Partition Template**.

    b) If you choose **Specify Virtual Networks in this Partition Template**, you must specify the virtual local area network (VLAN) Name and VLAN ID. To add a virtual network, click the **Add Network** tab. A row is appended to the bottom of the table with the appropriate fields. To remove a network, select a network to be deleted from the table and click **Remove Selected**.

    c) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

13. To change the virtual Network Interface Controller (vNIC) settings of the template, click the **Virtual NICs** tab in the **Details** area.

a) You can edit, add, or remove vNICs only when you clear the **Use Captured I/O Information** check box. To add vNICs to the template, click **Add Virtual NIC**. You can specify values for the **Port VLAN ID**, **VLAN ID restrictions**, **MAC Address**, **OS MAC Address restrictions**, **Backing Device Capacity (%)**, and **Backing Device Failover Priority** fields. The default value for the **Backing Device Failover Priority** field is Use Default. To modify the default values for the backing device or to add and remove backing devices to the vNIC, select the virtual NIC and click **Action** > **Modify Backing Device**. You can enable or disable Virtual NIC Auto Priority Failover. To add more backing devices to the vNIC, click **Add Backing Device**. To remove a backing device, select the backing device and click **Remove**. To change the settings of a vNIC, select the vNIC that you want to change and click **Actions** > **Modify**. To remove a vNIC, select a vNIC to be deleted from the table and click **Actions** > **Remove**.

   If the **Use Captured I/O Information** check box is selected, you can view the vNICs listed in the table but you cannot edit the vNICs. To view the details of a vNIC, select the vNIC and click **Actions** > **View**.

b) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

14. To change the virtual storage settings of the template, click the **Virtual Storage** tab in the **Details** area.

   a) In the **Virtual SCSI** tab, you can configure the virtual SCSI adapters that are required for a partition activation. If you are using a captured template, the table displays all the shared storage pool volumes that are captured in the template.

   b) In the **Shared Storage Pool Volume** area, click the **Add SSP Volume** tab to add shared storage pool volumes, or click **Remove Selected** to remove the selected shared storage volume.

   You must specify a shared storage pool volume in the template to add any shared storage pool volumes when you are creating the partition by using the template.

   c) You can choose to specify the shared storage pool cluster when you are creating the partition by selecting the value **Choose at deploy** from the `Shared Storage Pool Cluster Name` list, or you can specify the shared storage pool and the tier.

   d) You can also enable or disable thin provisioning for the shared storage volume pools.

   e) In the **Physical Volumes** area, you can enable **Configure Physical Volumes**.

   a) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

15. To change the virtual Fibre Channel settings of the template, click the **Virtual Fibre Channel** tab.

   a) You can choose **Configure Virtual Fibre Channel storage during deployment**, **Configure Virtual Fibre Channel storage with captured information**, or **Do not configure Virtual Fibre Channel storage**.

   If you are using a captured template, the captured Fibre Channel Port information is also displayed.

   b) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

16. To change the virtual optical device settings of the template, click the **Virtual Optical Device** tab. You can choose to configure the virtual optical device adapter that is required for the partition activation. The table displays all the Virtual Optical Devices that are captured in the template.

   a) Click **Add Virtual Optical Device** to add an optical device.

   b) To delete a device, click the **Remove** tab that is displayed in the row of the optical device that you want to delete.

   c) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

17. In the **Details** area, click the **Hardware Virtualized I/O** tab.

   a) Click the **HEA** tab. You can change the operating system-level VLAN ID and MAC address restrictions for each logical Host Ethernet adapter (LHEA) that is listed in the table. If you select the **Use Captured I/O Information** check box, you cannot add or remove any logical ports or change the values of the **Advanced Settings** tab.

You cannot add or remove any LHEA.

b) Click the **SR-IOV** tab. The **Use Captured I/O Information** check box is selected by default.

- If you are using captured I/O information, you cannot add or remove any logical ports or change the values in the **Advanced Settings** tab. If the captured partition template contains information about RDMA over Converged Ethernet (RoCE) logical ports, that information about the RoCE logical ports is also listed. In the **SR-IOV Logical Ports** area, you can view but cannot change the settings of the SR-IOV backup devices that can be migrated to another server.
- If you are not using the captured I/O information by clearing the **Use Captured I/O Information** check box, you can complete the following tasks:
  - Change properties that are specific to Ethernet logical ports.
  - Specify the capacity of the logical port as a percentage of the capability of the physical port. The capacity level determines the amount of resources that are assigned to the logical port from the physical port.
  - Click the **Add** tab to add a RoCE logical port to the template.
  - Configure an SR-IOV logical port such that you can migrate the logical port to another server by selecting the SR-IOV backup device from the **Backup Device Type** list in the **Modify Advanced Settings** page.

Notes:

- You can migrate only SR-IOV logical ports but you cannot migrate RoCE logical ports.
- When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

c) To edit the settings of the logical port, select the port and click the **Advanced Settings** tab. You can select values for the **OS MAC Address Restrictions** and **VLAN ID Restrictions** fields.

- If you selected **Allow Specified** as the value for the **OS MAC Address Restrictions** field, you must specify the MAC addresses in the **Specify allowed MAC Address(es)** field. To add more MAC addresses, click the plus sign (+) and to remove MAC addresses, click the minus sign (–).
- If you selected **Allow Specified** as the value for the **VLAN ID Restrictions** field, you must specify the VLAN ID or range of VLAN IDs in the **Specify VLAN ID(s) or range** field.
- If you specify the value for the VLAN ID as 0, the **802.1Q Priority** priority field is disabled. However, if you specify any value in the range 2 - 4094, you can set the priority value. Priority is used to prioritize the frames in a VLAN network.
- The **Promiscuous** option is disabled unless the logical port is used as the physical device for bridging virtual Ethernet adapters on client partitions. When the logical port is in promiscuous mode, the **VLAN ID Restrictions** and the **OS MAC Address Restrictions** fields are disabled. Click **Close**.

d) To add a logical port, click the **Yes** tab.

e) To remove a logical port, click the **Remove Selected** tab.

f) Click **Save As** to save the changes in a new name template. Otherwise, click **Save and Exit** to overwrite the changes in the template.

Changing a partition template to disable Live Partition Mobility

You can disable the Live Partition Mobility feature of a logical partition by changing a user-defined or captured partition template and save the changes in a new partition template by using the Hardware Management Console (HMC).

About this task

The HMC provides the **Disable Migration** option to disable the Live Partition Mobility feature at a logical partition level. This option can be used by customers to address application licensing requirements of Independent Software Vendors (ISV). To disable the Live Partition Mobility feature for a logical partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab and select the partition template that you want to change.
4. Click **Action** > **Edit**.
5. To change the partition properties of the template, click the **Properties** > **Advanced settings** tab.
6. Select the **Disable Migration** check box.
7. Click **Save and Exit** to save the changes in the same template, or select **Save As** to save the changes in a new template.

   Some Independent Software Vendors might require you to purchase a license for all systems where their application can be migrated. Rather than disabling the Live Partition Mobility feature at a system level, IBM provides this logical partition level mechanism that can be audited, to disable migration to address ISV licensing requirements while preserving the ability to leverage migration for applications running on other logical partitions on the system.

   Note: IBM software does not stipulate such licensing requirements.

   The **Disable Migration** option is supported on all firmware versions, and when the system is managed by an HMC that is at version 8.4.0, or later. Also, you can run the **chsyscfg** command with a value 1 for the *migration_disabled* attribute, from the HMC command line. To disable the Live Partition Mobility feature of a logical partition during partition creation, run the **mksyscfg** command with a value 1 for the *migration_disabled* attribute, from the HMC command line. The **Disable Migration** option is also supported by the Representational State Transfer (REST) application programming interfaces (APIs).

   Note: PowerVM NovaLink supports the **Disable Migration** option when the system is co-managed by an HMC. However PowerVM NovaLink does not provide an option to disable Live Partition Mobility feature.

Viewing system event logs for the Live Partition Mobility disable operation

Any changes that are made to the **Disable Migration** option provided by the Hardware Management Console (HMC) is logged as a system event, and can be checked for auditing purposes. A system event is also logged when the remote restart or simplified remote restart capability is set. The system event logs are read only and cannot be modified.

A system event is logged when the following actions occur:

- The remote restart, simplified remote restart, or Live Partition Mobility attributes are set during creation of a logical partition.
- The remote restart, simplified remote restart, or Live Partition Mobility attributes are changed.
- When you restore profile data. For more information about restoring profile data, see Restoring profile data.

You can view the system events by running the **lssvcevents** command from the HMC command-line interface. You can also view the system events by using the graphical user interface (GUI). For more information about using the GUI, see Console Events Logs. By running the **chhmc** command from the HMC command-line interface, these system events can also be sent to a remote server that is on the same network as the HMC.

The following system events can be logged:

Table 1. Event ID and the corresponding message string

Event ID	Event message String
2420	User name {0}: Disabled partition migration for partition {1} with ID {2} on managed system {3} with MTMS{4}
2421	User name {0}: Enabled partition migration for partition {1} with ID {2} on managed system {3} with MTMS{4}
2422	User name {0}: Disabled Simplified Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}
2423	User name {0}: Enabled Simplified Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}
2424	User name {0}: Disabled Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}
2425	User name {0}: Enabled Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}

The following are examples of System Events:

- Command to check which logical partitions managed by an HMC have the Live Partition Mobility feature disabled or enabled:

  ```
  lssvcevents -t console | grep vclient
  ```

  The following examples show a sample output of the **lssvcevents** command:

  ```
  time=10/30/2015 10:11:32,text=HSCE2521 UserName hscroot: Enabled partition migration for
  partition
  vclient10 with Id 10 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

  ```
  time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for
  partition
  vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

- Command to check which logical partitions managed by an HMC have the Live Partition Mobility feature disabled:

  ```
  lssvcevents -t console | grep HSCE2520
  ```

  The following example shows a sample output of the **lssvcevents** command:

  ```
  time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for
  partition
  vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

- Command to check which logical partitions managed by an HMC have the Live Partition Mobility feature disabled or enabled for a particular system (1234567):

  ```
  lssvcevents -t console | grep  "partition migration for partition" | grep 1234567
  ```

The following examples show a sample output of the **lssvcevents** command:

```
time=10/30/2015 10:11:32,text=HSCE2521 UserName hscroot:  Enabled partition migration for 
partition 
vclient10 with Id 10 on Managed system ct05 with MTMS 8205-E6D*1234567.
```

```
time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for 
partition 
vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567.
```

- Command to check whether a specific logical partition (vclient9) in a specific system (1234567) managed by an HMC has the Live Partition Mobility feature disabled or enabled:

```
lssvcevents -t console | grep  "partition migration for partition vclient9" | grep 1234567
```

The following example shows a sample output of the **lssvcevents** command:

```
time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for 
partition 
vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567
```

Copying a partition template

You can copy a quick-start or captured partition template into a new partition template along with the configuration details that are specified in the template by using the Hardware Management Console (HMC).

About this task

To copy a partition template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab and select the partition template that you want to copy and click **Action** > **Copy**.
4. In the **Copy Partition Template** page, specify the name for the template in the **Template name** field. If a template with that name exists, the copy fails and an error message is displayed.
5. Click **OK**.

Importing a partition template

You can import a partition template into the template library by using the Hardware Management Console (HMC).

About this task

Before importing a partition template, consider the following restrictions:

- If the partition template schema is different from the schema that is supported by the HMC, for example, if a tag that is not a part of the template OpenDocument Spreadsheets (ODS) file element is used, the partition template cannot be imported. However, if you are using HMC V9.1.930, or later, you can select another partition template and import that partition template.
- If the partition template file size exceeds 10 MB, the partition template cannot be imported and the operation fails. However, if you are using HMC V9.1.930, or later, you can select another partition template and import that partition template.

To import a partition template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab and click **Import**.
4. In the **Import Partition Template** page, click **Browse** to navigate to the template file.

   After you select the file, the selected file name is displayed in the **Template name** field. You can optionally change the name of the file. If a template with that name exists, the import fails and an error message is displayed. If you are using HMC V9.1.930, or later, you can change the name of the file or select another partition template and import that partition template.
5. Click **OK**.

Exporting a partition template

You can export a partition template from the template library by using the Hardware Management Console (HMC).

About this task

To export a partition template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab. Select the template and click **Action** > **Export**.

   A browser-generated window opens where you can choose to save the exported file.
4. Click **Save file** tab and specify the name of the file.
5. Click **OK**.

Deleting a partition template

You can delete a partition template from the template library by using the Hardware Management Console (HMC).

About this task

To delete a partition template, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Templates and OS Images**.
3. Click the **Partition** tab. Select the template and click **Action** > **Delete**.
4. In the **Delete Template** page, click **Yes** to delete the selected template. Otherwise click **No** to close the **Delete Template** page.

Creating a logical partition by using a template

You can create a partition by using partition templates that are available in the template library of the Hardware Management Console (HMC). The **Create a Partition from Template** wizard guides you through the deployment process and configuration steps.

About this task

The HMC checks whether the template that you selected matches the system capabilities when you click **Next**. If the template does not match the system capabilities, an error message is displayed. You can choose another template that matches the capabilities, or edit the template and use the changed template for creating the logical partition.

To create a partition by using the partition template, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
   a) Click **All Systems**. The All Systems page is displayed.
   b) In the work pane, select the system and click **Actions** > **View System Properties**. The Properties page is displayed.
   c) Expand **System Actions** > **Templates** > **Create Partition from Template**. Alternatively, you can create a partition by accessing the template library.
2. Click **Next**.

   If the selected template is compatible with the target system, the **Partition Configuration Summary** page is displayed. For servers that do not support IBM i partitions with native I/O capability, you must enable IBM i restricted I/O mode by selecting the **Restricted I/O Partition** check box. If you continue with the partition creation without selecting the **Restricted I/O Partition** check box, a warning message is displayed. You must run the wizard again and select the **Restricted I/O Partition** checkbox to continue with the partition creation.
3. In the **Partition Configuration Summary** page, you can change the default partition name. For AIX or Linux partitions, you can also select the **Shared Processor Pool** option if the partition template specifies that the partition uses shared processors. For IBM i logical partitions, the **IBM i Tagged I/O** tab is displayed (for more information on **IBM i Tagged I/O** tab, see step 8). Additionally, if the managed system supports virtual serial number (VSN) and the managed system is not in a Power® Enterprise pool 2.0, you can also specify the **Virtual Serial Number** for the logical partition. You can choose any one of the following options:

   - No VSN
   - Auto-assign
   - Select from the pool

   Click **Template Details** to view the details of the template. If you are using HMC V9.1.920, or later, and when the firmware is at level FW920, or later, the secure boot feature is supported. To select a value for the **Secure Boot** field, click the **General** tab, and click **Advanced**. Select a value and click **Next**. Additionally, if you are using HMC V9.2.950, or later, and when the firmware is at level FW950, or later, you can specify a value of either 0 kilobytes (KB) or a value that is supported by the managed system for the **Keystore Size** field. To select a value for the **Keystore Size** field, click the **General** tab, and click **Advanced**.

   Notes:

   - If the system key for encrypting the partition keystore data is not a user-defined key, then you cannot create an logical partition with both the partition keystore and simplified remote restart capabilities enabled.

- If the template contains details about the virtualization capabilities like Live Partition Mobility, you can view the details by clicking **Template Details**.

4. In the **Persistent Memory** page, complete the following steps:

   If you are using a captured partition template, the **Persistent Memory** area shows all the persistent memory volumes that are captured from an original system. You can edit the name for the persistent memory volume and associate the persistent memory volumes to the logical partition. If you are not using the captured I/O information, you can add a persistent memory volume to the logical partition by using the using the **edit partition template** task. During the deploy operation, you can change the name and size of the persistent memory volume and associate it with the logical partition.

5. In the **Physical I/O** page, complete the following steps:

   a) In the **Physical I/O Adapters** area, you can select the physical I/O adapters for the logical partition. To view the adapters that are available in other drawers of the system, select the drawer from the **View adapters in** list.

      - If you are using V9.1.910, or earlier, and if you are using a captured partition template and if the template contains captured I/O information, the physical I/O adapter information that is captured from the partition during the capture operation is not used by the HMC. The HMC checks for the available physical I/O adapters on the destination server that can be assigned to the logical partition. These available physical I/O adapters are listed in the **Physical I/O** page, and you can choose a physical I/O adapter for assignment to the logical partition from this list. This scenario is also applicable when you do not want to use the captured I/O information (not selecting the **Use Captured I/O Information** checkbox in the **Physical I/O** page).
      - If you are using HMC V9.1.930, or later, and if you are using a captured partition template, the physical I/O adapters on the destination server that match the details of the physical I/O adapters in the captured template are selected automatically. If there are no physical I/O adapters that match the details of the captured template, the HMC checks for the available physical I/O adapters on the destination server that can be assigned to the logical partition. The available physical I/O adapters are listed in the **Physical I/O** page, and you can choose a physical I/O adapter for assignment to the logical partition from this list.

   b) In the **SR-IOV Logical Ports** area, you can complete the following tasks:

      - If you are using the captured I/O information and the captured partition template contains information about RDMA over Converged Ethernet (RoCE) logical ports, information about the RoCE logical ports is listed. You can select the target port from the list of target physical ports when there are multiple target physical ports that match the location code and support the logical port protocol of the source logical port (Ethernet or RoCE).

        Notes:

        – If you are using HMC V9.1.920, or earlier, the deployment operation fails when the server does not contain the exact match of the required SR-IOV Logical Ports.

        – If you are using HMC V9.1.930, or later, you can continue with the deployment operation although the server does not have an exact match of the required SR-IOV logical ports, but has alternative SR-IOV physical ports to create logical ports. Also, you can continue with the deployment operation if the server does not have an exact match of the required SR-IOV Logical Ports for the backing devices, but has alternative SR-IOV physical ports to create logical ports that can be configured as backing devices.

        – You can view but cannot change the settings of the SR-IOV logical ports that can be migrated to another server.
      - If you are not using the captured I/O information from the template, you can complete the following tasks:

        – Assign the SR-IOV logical ports that are displayed to the logical partition. You can also specify the advanced properties for the logical port.

        – Select the target port, and change the capacity of the logical port. Only the physical ports that support the logical port protocol (Ethernet or RoCE) are listed.

- Configure an SR-IOV logical port such that you can migrate the logical port to another server.

Notes:

- You can migrate only SR-IOV logical ports but cannot migrate RoCE logical ports.
- When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a technology preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

c) In the **Configure Backup Device Type** area, you can configure the **Backup Device Type** as a Virtual NIC adapter or a Virtual Ethernet Adapter.

6. If you specified the networks in the template before creating the partition, the **Network Configuration** page displays a summary of the virtual networks to which the partition connects. The page displays a list of available networks if you did not specify the networks before you created the partition. In both cases, you can specify the virtual Ethernet adapter ID. Click **Next**.
7. In the **Virtual NIC Configuration** page, the table displays all the virtual Network Interface Controllers (vNICs) that are in the template. When the template uses captured I/O information, you can change the value of the **Hosting Partition** field and select a physical port that matches the hosting partition. When the template does not use captured I/O information, you can change the value of the **Capacity (%)**, **Physical Port**, **Hosting Partition**, and **Failover Priority** fields for each backing device. The number of physical ports must be equal to or greater than the number of backing devices of the vNICs that are specified in the template because each backing device of a vNIC must be created on a different physical port. You can also specify values for the **MAC Address** field which depends on the template settings. Click **Next**.
8. In the **Storage Configuration** page, complete the following steps:

   Note: You can configure storage resources such as virtual Small Computer Serial Interface (SCSI), virtual Fibre Channel, and virtual Optical devices. When you are using virtual adapters, you can connect logical partitions with each other without using physical hardware. Operating systems can display, configure, and use virtual adapters similar to how it can display, configure, and use physical adapters. Depending on the operating environment used by the logical partition, you can create virtual Ethernet adapters, virtual Fibre Channel adapters, virtual Optical devices, and virtual SCSI adapters. You can use virtual SCSI (vSCSI) to simplify the backup and maintenance operations on your managed system. When you back up the data on the system logical partition, you also back up the data on each client logical partition. You can configure the managed system with N_Port ID Virtualization (NPIV), such that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter. NPIV is a standard technology for Fibre Channel networks. NPIV enables you to connect multiple logical partitions to one physical port of a physical Fibre Channel adapter.

   a) Click **Virtual SCSI**.

   b) In the **Physical Volume** area, you can assign physical volumes. Click **Edit Connections** to edit the Virtual I/O Server (VIOS) connections for the physical volumes. Click **Show assigned physical volumes** to view more physical volumes in the table.

   c) In the **Shared Storage Pool Volumes** area, you can view the device details for a shared storage pool volume. Click the device name to open the **Configure SSP Volume** window. In the **Configure SSP Volume** window, select the shared storage pool cluster and the tier that you want to assign the device to.

   You can also enable or disable thin provisioning for the device. In a thin-provisioned device, the used storage space might be greater than the actual used storage space. If the blocks of storage space in a thin-provisioned device are unused, the device is not entirely backed up by physical storage space. By using thin-provisioning, you can exceed the storage capacity of the storage pool.

   d) Click **Virtual Fibre Channel**. The **Virtual Fibre Channel** content area displays a table with the available Virtual Fibre Channel ports to which the partition can connect. You can select the Virtual Fibre Channel port from the displayed list of ports.

  - If you are using HMC V9.1.920, or earlier, and if you are using a captured partition template, the deployment operation fails when the server does not contain the exact match for the required Fibre Channel adapters.
  - If you are using HMC V9.1.930, or later, and if you are using a captured partition template, the Virtual Fibre Channel adapters on the destination server that match the details of the Virtual Fibre Channel adapters in the captured template are selected automatically. When there are no Virtual Fibre Channel adapters that match the details of the captured template, the HMC checks for the available Virtual Fibre Channel adapters on the destination server that can be assigned to the logical partition. The available Virtual Fibre Channel adapters are listed in the **Virtual Fibre Channel** page, and you can choose a Virtual Fibre Channel adapter for assignment to the logical partition from this list.

e) Click **Virtual Optical Devices**.

The **Virtual Optical Devices** content area displays the devices that are specified in the template. You can optionally change the VIOS on which the device is created.

f) Click **Next**.

9. For IBM i logical partitions, the **IBM i Tagged I/O** page is displayed. Select values for the **Load Source**, **Alternate Restart Device**, **Console**, **Alternate Console**, and **Operations Console** fields.
10. In the **Summary page** page, you can view a summary of the changes. Select one of the following options:

    - **Activate partition** - Creates the partition with the resources you selected in this wizard and activates the partition.
    - **Create partition and Apply Configuration** - Creates the partition with resources that you selected in this wizard.

    If you chose this option in step 9, the logical partition is created with the resources that you selected. This step might take some time. You can view the progress of the operation. After the operation completes, select **Click to Install** if you want to install the operating system on the logical partition by using the **Network Boot** wizard.
11. Click **Finish**.

    Note: If you are using HMC V9.1.920, or earlier, the deployment operation fails and a partition is not created when the system does not match all the configuration along with the I/O adapter configuration that is provided by the selected template. If the virtual network or virtual storage mapping fails, the partition is created and a message is displayed that indicates that the partition was created with errors. The created partition cannot be rolled back. You must manually delete the partition or go to the **Manage PowerVM** functions available in the HMC to assign the network or storage to the created partition.

Creating logical partitions by using Create partition option

You can create an AIX, Linux, or IBM i logical partition by using the **Create Partition** option.

About this task

To create an AIX, Linux, or IBM i logical partition by using the **Create Partition** option, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The All Systems page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Partitions**. The **Partitions** page is displayed.

4. Click **Actions** > **Partitions**.
5. Click **Create Partition**.

   By default the logical partition is created in the shared mode with 0.1 processing units for the **Maximum**, **Allocated**, and **Minimum** fields and one virtual processor for the **Maximum**, **Allocated**, and **Minimum** fields. When the server does not support shared processor pools, the logical partition is created in the dedicated mode with one processor for the **Maximum**, **Allocated**, and **Minimum** fields. The default value for the **Maximum** field is 4 GB, and 1 GB for the **Allocated**, and **Minimum** fields, for both the shared and dedicated modes. Before you activate the logical partition, you must assign storage and network resources to the logical partition. Optionally, you can change the default values that were assigned by using the Manage Partition functions.

   Alternatively, you can select the **Assign All System Resources** checkbox to assign all the resources to the partition. The logical partition creation succeeds only when all the other active logical partitions and Virtual I/O Servers are shut down.
6. To create a single logical partition, complete the following steps:
   a) Select the **Single Partitions** tab. Click the **Basic Configuration** tab to specify the **Partition Name**, the **Partition ID**, and the **Partition Type**. Specify the maximum number of virtual adapters for the logical partition in the **Maximum Virtual Adapters** field. Additionally, if the managed system supports virtual serial number and the managed system is not in a Power Enterprise pool 2.0, you can also specify the **Virtual Serial Number** for the logical partition. You can choose any one of the following options:

      - No VSN
      - Auto-assign
      - Select from the pool

      Choose the **Select from the pool** option, to open the **Virtual Serial Number Selection** window. The window lists the virtual serial number groups and the available virtual serial number that can be assigned to the partition. Select a virtual serial number from the list.

      Note: When the Power Firmware Level is at FW950 and when the managed system has logical partitions that are assigned with virtual serial numbers, the managed system cannot be added to a Power Enterprise pool 2.0. Alternatively, if the managed system exists in a Power Enterprise pool 2.0, the managed system cannot assign a virtual serial number to the logical partition.

   b) Click the **Processor Configuration** tab to specify whether the logical partition uses processors that are dedicated to the logical partition, or processors that are shared with other logical partitions.

      - When you choose to create a logical partition with dedicated processors, you can specify values for the **Minimum**, **Allocated**, and **Maximum** fields.
      - When you choose to create a logical partition with shared processors, you can specify values for the **Minimum**, **Allocated**, and **Maximum** fields in the **Virtual Processors** area, and the **Minimum**, **Allocated**, and **Maximum** fields in the **Processing Units** area.
      - You can set the processor weight as **Capped** or **Uncapped**. When you set the processor weight as **Uncapped**, you must specify a value for the processor weight in the **Weight** field.

   c) Click the **Memory Configuration** tab to specify the values for the **Minimum**, **Allocated**, and **Maximum** fields.
7. To create multiple logical partitions, complete the following steps:
   a) Select the **Multiple Partitions** tab.
   b) To create a new logical partition with minimum resources that are assigned to the logical partition, click the **+** tab. The **Create Partition(s)** page lists the logical partitions with details about the logical partitions such as the **Partition Name**, **No. Of Instances**, **Partition ID**, **Partition Type**, **Maximum Virtual Adapters**, **Desired Memory (MB)**, **Processor Type**, and **Desired Processor**. Additionally, if the managed system supports virtual serial number and the managed system is not in a Power Enterprise pool 2.0, you can also specify the **Virtual Serial Number** for the logical partition. You can choose any one of the following options:

- No VSN
- Auto-assign

Note: When the Power Firmware Level is at FW950 and when the managed system has logical partitions that are assigned with virtual serial numbers, the managed system cannot be added to a Power Enterprise pool 2.0. Alternatively, if the managed system already exists in a Power Enterprise pool 2.0, the managed system cannot assign a virtual serial number to the logical partition.

c) To specify the processor and memory resources for a logical partition, select the logical partition from the **Partition Name** column.

d) Click the **Processor Configuration** tab to specify whether the logical partition uses processors that are dedicated to the logical partition, or processors that are shared with other logical partitions.

- When you choose to create a logical partition with dedicated processors, you can specify values for the **Minimum**, **Allocated**, and **Maximum** fields.
- When you choose to create a logical partition with shared processors, you can specify values for the **Minimum**, **Allocated**, and **Maximum** fields in the **Virtual Processors** area, and the **Minimum**, **Allocated**, and **Maximum** fields in the **Processing Units** area.
- You can set the processor weight as **Capped** or **Uncapped**. When you set the processor weight as **Uncapped**, you must specify a value for the processor weight in the **Weight** field.

e) Click the **Memory Configuration** tab. You can specify the minimum, allocated, and maximum amounts of memory resources that you want for the logical partition by specifying these values in the **Minimum**, **Allocated**, and **Maximum** fields.

f) Specify the maximum number of virtual adapters for the logical partition in the **Maximum Virtual Adapters** field.

g) Specify the number of instances for the logical partition in the **No. Of Instances** field. The maximum value that you can specify for this field is 20.

h) To delete a logical partition, select the logical partition from the **Partition Name** column and click the **-** tab.

8. Click **OK**.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled “Cookies, Web Beacons and Other Technologies”.

Programming interface information

Setting up the virtualization environment documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.2, IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.4, and IBM Virtual I/O Server Version 3.1.2.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Managing the virtualization environment

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 109.

This edition applies to IBM® AIX® Version 7.2, to IBM AIX Version 7.1, to IBM AIX Version 6.1, to IBM i 7.4 (product number 5770-SS1), to IBM Virtual I/O Server Version 3.1.2, and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

Contents

Managing the virtualization environment

You can use the PowerVM® management, Virtual I/O Server management, and partition management functions available with the Hardware Management Console (HMC) Version 8, Release 8.1.0, Service Pack 1, or later to manage the virtualization capabilities of your IBM Power Systems servers.

What's new in Managing the virtualization environment

Read about new or changed information in Managing the virtualization environment since the previous update of this topic collection.

November 2020

The following information is a summary of the updates made to this topic collection:

- Multiple topics have been updated with information about enhancements to the HMC graphical user interface.
- Added information about validating the VIOS for maintenance readiness in the topic "Validating the Virtual I/O Server for maintenance readiness" on page 4.
- Added information about the virtual serial number (VSN) in the topic "Changing partition properties and capabilities" on page 63.
- Added information about the platform keystore capability in the topic "Changing advanced partition settings" on page 68.

April 2020

- Updated the topic "Virtual network bridges" on page 15 with information about the network bridge configuration.

October 2019

- The following topics were added or updated with information about configuration, validation and migration of a logical partition that is configured with the SR-IOV logical ports:
  - "Validating the configuration of a logical partition before the migration operation" on page 71
  - "Migrating a logical partition" on page 72
  - "Changing partition profile properties" on page 82
  - "Adding SR-IOV logical ports" on page 100
- Added information about the persistent memory volume in the topic "Managing persistent memory volume" on page 77
- Updated the following topics with information about the capabilities of a restricted I/O partition:
  - "Changing partition properties and capabilities" on page 63
  - "Changing advanced partition settings" on page 68
- Updated the topic "Viewing virtual Fibre Channel adapters" on page 32 with information about the capabilities of the Fibre Channel adapter.

May 2019

- The following topics were updated with information about enhancements to the HMC graphical user interface:
  - "Managing system properties" on page 3
  - "Managing the properties of a Virtual I/O Server" on page 9

  - “Viewing the virtual network configuration” on page 16
  - “Activating IBM i partitions” on page 60
  - “Activating AIX or Linux partitions” on page 61
- The following topics were added or updated with information about RDMA over Converged Ethernet (RoCE) support:
  - “Managing partition profiles for logical partitions” on page 80
  - “Creating a partition profile” on page 80
  - “Copying a partition profile” on page 81
  - “Changing partition profile properties” on page 82
  - “Deleting a partition profile” on page 84
  - “Adding SR-IOV logical ports” on page 100

August 2018

- The following topics were added or updated with information about enhancements to the HMC graphical user interface:
  - “Changing processor settings” on page 73
  - “Changing memory settings” on page 75
  - “Managing physical I/O adapters” on page 78
  - “Managing virtual SCSI resources for a partition” on page 93
  - “Synchronizing a virtual switch” on page 22
- The following topics were updated with information about the secure boot capability:
  - “Changing advanced partition settings” on page 68
  - “Managing the properties of a Virtual I/O Server” on page 9
- Added information about Universal Serial Bus (USB) flash drive support for VIOS installation in the “Activating Virtual I/O Servers” on page 6 topic.

Managing systems

You can use the PowerVM function on the Hardware Management Console (HMC) Version 8, Release 8.1.0, Service Pack 1, or later to manage the system-level virtualization capabilities of IBM Power Systems, such as managing a Virtual I/O Server (VIOS), managing virtual networks, managing virtual Network Interface Controllers (vNICs), and managing virtual storage.

If you are using an HMC interface, you can perform system management functions, such as configuring a Virtual I/O Server (VIOS), virtual networks, and virtual storage, by accessing the options listed under the PowerVM area of the graphical user interface.

You can manage the system-level virtualization capabilities of IBM Power Servers only when a server is managed by the HMC, or when a server is co-managed by the HMC and PowerVM NovaLink, with the HMC or PowerVM NovaLink in the master mode. The PowerVM NovaLink architecture enables management of highly scalable cloud deployment by using the PowerVM technology and OpenStack solutions. The architecture provides a direct OpenStack connection to a PowerVM server. The NovaLink partition runs the Linux operating system and the partition runs on a server that is virtualized by PowerVM. The server is managed by PowerVC or other OpenStack solutions.

If you want to manage system-level virtualization capabilities by using the HMC, you must set the HMC or the PowerVM NovaLink to the master mode. Run the following command from the command line to switch the HMC to the master mode:

```
chcomgmt -m <managed system> -o setmaster -t norm
```

Managing system properties

You can view and change the properties of the selected managed system. You can view the capabilities that are supported by the managed system.

Before you begin

Stealable processor or memory value indicates the processor and memory resources that are retrieved by the HMC from shutdown or hibernated partitions. You can view information about the estimated available resources including, stealable processor and memory resources, which can be used to perform certain partition management functions. You can also view information about the estimated remaining processor resources for AIX or IBM i partitions that are in the running state.

About this task

To view and change the properties of the selected managed system, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the navigation pane, ensure that **Properties** is expanded.
   a) Click **General Settings** > **General Properties**. You can view and change the general system properties. You can change the system's name, location, and description, assigned service partition (if designated), power-off setting, and group tags. You can view only the reference code, machine type, serial number, managed system firmware, default configuration, and the maximum number of partitions that can be defined on the server.

   b) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

   c) Click **General Settings** > **Migration**. You can view or change the partition mobility properties and change the migration policy for inactive partitions on the managed system.

   - Select the migration policy that you want to use when you migrate inactive partitions. You can select one of the following policies:
     - `Partition configuration`: Configures the management console to use the partition state that is defined for the logical partition when you migrate an inactive partition. If the inactive partition cannot start automatically, the management console uses the configuration data that is defined for the partition in the last activated profile.
     - `Last Activated Profile`: Configures the management console to use the memory and processor configuration data that is defined in the last activated profile for the partition when you migrate an inactive logical partition.
   - Select **Allow Migration with Inactive Source Storage VIOS** to perform Live Partition Mobility (LPM) when the source Virtual I/O Server (VIOS) that is hosting the storage adapters is powered off or shutdown. If you enable this feature, the storage configuration-related data is collected for all client partitions based on the CEC level preference. The collected data is used to perform LPM on the powered off VIOS.
   - View the migration capabilities table to see information about the type of migration that is supported, number of migrations in progress, and number of migrations that are supported by the managed system.

   d) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

e) Click **General Settings** > **Power-On Parameters**.

   You can change the power-on parameters for the next system restart by changing the values in the **Next Value** fields. The **Current Value** field displays the value that was used when the system was last restarted. You can change the value for the partition start policy, power-on side, keylock position, IBM i IPL source, and AIX/Linux boot mode. The changed value takes effect after the next system restart.

f) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

g) Click **General Settings** > **Advanced**.

   You can view or change the settings for Barrier Synchronization Register (BSR), huge page memory, processor performance, memory mirroring, memory optimization, Virtual Trusted Platform Module (VTPM) enabled partitions, and supported hardware accelerator types for the managed system. The Nutanix hardware accelerator has a total of 32 hardware accelerator types. Currently, only the GZIP co-processor type is supported. You can increase the amount of available mirrored memory on the system and perform defragmentation operation, by using the **Memory Optimization Tool**.

h) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

5. In the navigation pane, click **Processor, Memory, I/O** to view the memory, processor, and physical I/O resource settings for the managed system. You can click **I/O Pools** to display all the I/O pools available in the managed system. The maximum number of I/O pools that are allowed is 1000.
6. Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.
7. In the navigation pane, click **Hardware Virtualized I/O**. The Single Root I/O Virtualization (**SR-IOV**) page for the selected Virtual I/O Server is displayed in the work pane.

   a) The **SR-IOV** page lists all of the SR-IOV logical ports that are connected to the VIOS. Right-click a logical port and select **Modify Port** or **Remove Port** to change or remove the selected port. Click **Add Port** to add an SR-IOV logical port to the VIOS partition.

   b) The **HEA** page lists all of the Logical Host Ethernet Adapters (LHEAs) connected to the VIOS. Select an LHEA adapter from the list to view the port configuration details. Right-click any port in the table to modify the port configuration and view the partitions that are associated with the selected HEA port.

   c) In the **HCA** page, click **Launch Manage Host Channel Adapters** to open the HMC panel with a list of available HCAs. Select an HCA to display the current partition usage for the selected HCA.

Managing Virtual I/O Servers

You can manage a Virtual I/O Server (VIOS) by using the **Virtual I/O Servers** option listed under the PowerVM area of the interface available in the Hardware Management Console (HMC).

The **Virtual I/O Servers** option displays a list of Virtual I/O Servers that are configured in the managed system. It also displays information about each VIOS configuration such as allocated memory, allocated processing units, allocated virtual processors, RMC status property, operating system (OS) version information, and status.

Note:

- The suggested VIOS level is 2.2.3.3, or later. If the VIOS is not at the suggested level, you might not get optimal performance and certain functions such as Shared Storage Pool management are not available.
- If your VIOS license is not accepted, some of these properties are not populated and you are not able to fully manage the VIOS. When your VIOS license is not accepted, the OS version information shows the version as **`License not accepted`**.

Validating the Virtual I/O Server for maintenance readiness

When the Virtual I/O Server (VIOS) partition is in **Running** state with an active Resource Monitoring and Control (RMC) connection, and you have access to all the resources of the managed system, you can validate the VIOS for maintenance readiness by using the Hardware Management Console (HMC). You can view the impacted client partitions for storage or network redundancy on the resources provided by VIOS.

About this task

To validate the VIOS for maintenance readiness, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the server name where you want to activate the VIOS and click **Actions** > **View System Properties**.
4. In the **PowerVM** area, click **Virtual I/O Servers**. The Virtual I/O Servers that are available on the system are displayed.
5. In the work pane, select the VIOS and click **Actions** > **Validate Maintenance Readiness**. The **Validate Maintenance Readiness** window is displayed.
6. In the **Validate Maintenance Readiness** window, the following sections are displayed:

   - **All**: Select the **All** option to view both the errors and the warning message information related to storage or network redundancy. By default, the **All** option is selected.
   - **Errors**: Select the **Errors** option to view only the error message information related to storage or network redundancy.
   - **Warnings**: Select the **Warnings** option to view only the warning message information related to storage or network redundancy.

   a) **Virtual SCSI Storage Validation** - Click and expand the **Virtual SCSI Storage Validation** section. The **Virtual SCSI Storage Validation** section displays the following information:

      - **Partition Name (State)**: Displays the name and state of the partition.
      - **Storage Name**: Displays the name of the storage device.
      - **Storage Type**: Displays the type of the storage such as physical volume, logical volume, virtual optical media and logical units.
      - **Remarks**: Displays the errors and the warning message information related to storage redundancy.

   b) **Virtual Fibre Channel Validation** - Click and expand the **Virtual Fibre Channel Validation** section. The **Virtual Fibre Channel Validation** section displays the following information:

      - **Partition Name (State)**: Displays the name and state of the partition.
      - **VFC Host Adapter**: Displays the name of the virtual Fibre Channel host adapter.
      - **Remarks**: Displays the errors and the warning message information related to virtual Fibre Channel host redundancy.

   c) **Virtual NIC Validation** - Click and expand the **Virtual NIC Validation** section. The **Virtual NIC Validation** section displays the following information:

      - **Partition Name (State)**: Displays the name and state of the partition.
      - **VNIC Device**: Displays the virtual NIC adapter value.
      - **Remarks**: Displays the errors and the warning message information related to virtual NIC adapter redundancy.

   d) **Virtual LAN Validation** - Click and expand the **Virtual LAN Validation** section. The **Virtual LAN Validation** section displays the following information:

      - **Partition Name (State)**: Displays the name and state of the partition.
      - **Port VLAN ID**: Displays the Port VLAN ID value.

    - **Virtual Switch**: Displays the name of the virtual switch.
    - **Virtual Network Name**: Displays the name of the virtual network.
    - **Remarks**: Displays the errors and the warning message information related to virtual network redundancy.

7. Click **Re-Validate** in the upper-right corner of the **Validate Maintenance Readiness** window to validate again and view the impacted client partitions for storage or network redundancy.
8. Click **View System VIOS** in the upper-right corner of the **Validate Maintenance Readiness** window to view all the Virtual I/O Server information of the managed system. The **View System Virtual I/O Server Information** window is displayed. The **View System Virtual I/O Server Information** window displays the following information:
    - **Name(ID)**: Displays the name of the Virtual I/O Server.
    - **State**: Indicates the current state of the Virtual I/O Server.
    - **RMC State**: Indicates the status of the Resource Monitoring and Control (RMC) connection.
    - **Remarks**: Displays the errors and the warning message information.
    - Click **Close** to close the **View System Virtual I/O Server Information** window.
9. Click **Close** to close the **Validate Maintenance Readiness** window.

Activating Virtual I/O Servers

You can activate **Virtual I/O Servers** by using the Hardware Management Console (HMC).

About this task

To activate the Virtual I/O Server (VIOS) and to set the activation options to activate or network boot the (VIOS) by using the activation wizard, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the server name where you want to activate the VIOS and click **Actions** > **View System Properties**.
4. In the **PowerVM** area, click **Virtual I/O Servers**. The Virtual I/O Servers that are available on the system are displayed.
5. In the work pane, select the VIOS and click **Actions** > **Activate**. The **Activate <VIOS partition name>** wizard is displayed.
6. From the **Select VIOS Configuration** list, select the required partition configuration profile.

    You can select only the profile that is associated to the selected partition. When you create a partition a default profile is always associated to the partition. This is indicated with the profile name being followed by **default** in parentheses.

    Note: If you choose **Current Configuration**, the **Advanced Settings** are unavailable.
7. From the **Activation Options** list, select the activation option for the partition.
    - Select **Activate** to activate the partition.

      Note: If you select **Activate**, the **Next** button is not available and you can only click **Finish** to activate and close the wizard after you make all your choices in the wizard.
    - Select **Install** to install the operating system on the partition. The HMC enables the network installation. When you select Install, click **Next** to configure the network settings for the logical partition.

8. Click **Advanced Settings** if you want to view and modify the following options for the selected partition:

   - **Keylock Position** establishes the power-on and power-off modes that are allowed for the system. You can select the following keylock values - Do not override configuration, Manual (attended), and Normal (unattended).

     ⚠ **Attention:** The **Manual** (attended) value is not preferred value for security reasons.

   - **Boot Mode** indicates the activation type for a partition. This activation type is applicable only for AIX, Linux, or Virtual I/O Server partitions. This option is not displayed for IBM i partitions.
   - **Open vterm** opens a virtual terminal console.
   - **Use VSI Profile** activates the partition with Virtual Station Interface (VSI) profiles.

     Note: If the VSI attributes are not set correctly, the activation fails.

9. If you selected **Activate** from the **Activation Options** list, click **Finish** to activate the VIOS partition and close the activation wizard.
10. If you selected **Install** from the **Activation Options** list, click **Next** to configure the network settings for the VIOS partition and to install the VIOS software. The **VIOS Installation Configuration** page opens.
11. In the **VIOS Installation Configuration** page, select one of the following methods from the **Installation Method** list to install a VIOS software on the VIOS partition:

    - **NIM Server**. You must enter the NIM server IP address to be accessed by the HMC. The NIM server IP address is the HMC IP address from which the VIOS IP address can be accessed. Additionally, you can view the MAC address of the system.
    - **Management Console Image**. You must enter the HMC IP address. You must also select VIOS image from the list.
    - **Management Console Session**. You must specify the boot mode to start the operating system on the logical partition. The valid boot modes are - **Normal**, **System Management Services (SMS)**, and **Open Firmware OK**.
    - **USB Image**. You must enter the HMC IP address to specify which Ethernet adapter must be used for communication with the VIOS. The **VIOS Installation Image** field lists all the Universal Serial Bus (USB) flash drives that are attached to your system. Select the VIOS image that you have saved in the USB flash drive and proceed with the VIOS installation.

12. Click **Advanced Settings** to view and change the following network configuration settings for the selected partition:

    a) From the **Adapter Speed** list, select the speed of the Ethernet adapter for the target partition. By default, **Auto** is selected to enable the system to determine the required speed for the adapter. You can also select the following values - **10**, **100**, or **1000**.

    b) From the **Adapter Duplex** list, select duplex value for the Ethernet adapter. By default, **Auto** is selected to enable the system to determine the required duplex for the adapter. You can also select the **Full** or **Half** values.

13. Click **Next**. The **VIOS Installation Progress** page is displayed.
14. In the **VIOS Installation Progress** page, you can install and activate the VIOS software and activate a VIOS partition on the managed system.
15. Click **Start** to start the VIOS software installation on the VIOS partition and then accept the licenses for each VIOS.
16. Click **Finish** to complete the VIOS software installation and close the activation wizard.

Viewing the configuration details of a Virtual I/O Server

You can view the configuration details of Virtual I/O Server (VIOS) resources on a system that is managed by a Hardware Management Console (HMC).

Procedure

To view resource information for a VIOS, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, click the server name that has the VIOS.
4. In the **PowerVM** area, click **Virtual I/O Servers**. The Virtual I/O Servers that are available on the system are displayed.
5. Select the VIOS and click **Actions** > **View Virtual I/O Server Properties** You can view the details of the VIOS configuration.

Adding a Virtual I/O Server

You can add one or more Virtual I/O Servers and configure virtual resources by using the **Add Virtual I/O Server** wizard in the Hardware Management Console (HMC).

About this task

You can add a Virtual I/O Server (VIOS) to provision system resources virtually to client partitions. Adding more Virtual I/O Servers can increase resource availability.

Procedure

To add a VIOS by using the **Create Virtual I/O Server** wizard, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the server name where you want to add the VIOS.
4. In the **PowerVM** area, click **Virtual I/O Servers**. The Virtual I/O Servers that are available on the system are displayed.
5. In the work pane, click **Create Virtual I/O Server**. The **Add VIOS Wizard** opens and displays the **General** tab.
6. Specify a name and a partition ID for the VIOS partition.
7. Click **Next**.
8. In the **Processor** tab, select the processor mode and change the maximum, allocated, and minimum processor resources that are assigned to the partition.
9. In the **Processor** > **Advanced Settings** section, select the appropriate setting for **Idle Processor Sharing**.
10. Click **Next**.
11. In the **Memory** tab, select the dedicated memory properties for the VIOS. You can change the maximum, allocated, and minimum memory values that are to be assigned to the VIOS.
12. Click **Next**.
13. In the **Physical I/O** page, assign the physical I/O adapters and Host Ethernet Adapters (HEA) to the VIOS.

Note: You must select one or more I/O adapters that can provide network and storage connectivity for the VIOS. Otherwise, the VIOS partition is created, but the VIOS installation and the VIOS deployment process fails.

14. Click **Next**.
15. In the **Configuration Summary** page, review the summary of configuration for the new VIOS. Select one of the following options to add the VIOS to the managed system:
    - **Apply configuration**: Creates the VIOS with the resources that you selected in this wizard. When you select this option, all the VIOS configurations are saved in the hypervisor, and the created VIOS is not powered on.
    - **Create Virtual I/O Server and Install Image**: Creates the VIOS by installing the VIOS image. When you select this option, you are directed to the **Install VIOS Wizard** where additional installation steps must be performed. In the **Install VIOS Wizard**, you can install the VIOS software on the VIOS partition that is created by using different installation methods. You can also provide network settings and accept the VIOS license by using this wizard.
16. Click **Finish** to create the VIOS on the managed system.

Managing the properties of a Virtual I/O Server

You can view, remove, or change the resources that are allocated to a Virtual I/O Server (VIOS) by using the **PowerVM** function in the Hardware Management Console.

About this task

You can change the resources that are configured for a VIOS.

Note: You can change only certain attributes while the VIOS is in active state. You can change all of the VIOS attributes when it is in inactive state.

Procedure

To view and change resources and the configuration for a VIOS, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the server name that has the VIOS that you want to change.
4. In the **PowerVM** area, click **Virtual I/O Servers**. The Virtual I/O Servers that are available on the system are displayed.
5. In the work pane, select the VIOS for which you want to view and change properties.
6. Click **Actions** > **View Virtual I/O Server Properties**.
7. In the navigation pane, ensure that **Properties** is expanded and **General Properties** is selected.

   a) In the **General Properties** page, you can view or change the VIOS name, VIOS version, IP address, boot mode, resource configuration, key lock position, view the machine serial number and machine type, and change the description and group tags.

   Note: If the Resource Monitoring and Control (RMC) connection is active, and if you want to accept the VIOS license, click **License not accepted** in the **VIOS Version** field to accept the VIOS license. The **Accept License for VIOS** window is displayed. Click **Accept** to accept the VIOS license. Alternatively, click **Cancel** to reject the changes and to close the window.

   b) In the **General Properties** page, click **Advanced** to enable or disable automatic start of the managed system, to enable or disable the mover service partition (MSP), to enable connection monitoring, to enable redundant error path reporting, to enable time reference, to enable VTPM, to allow performance information collection, to enable or disable the secure boot feature by selecting a value from the **Secure Boot** list, to specify the GZIP Quality of Service (QoS) credits, or to select the processor compatibility mode for a partition. You can also save the current configuration of the VIOS to a new partition profile, if you are a super administrator, service representative, operator, or product engineer. Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

8. In the navigation pane, click **Processors**.

   a) In the **Processors** page, select the values for virtual processors and values for processing units for the VIOS. You can set the VIOS to be either capped or uncapped. Click **Advanced** to select the processor compatibility mode and choose when to share a processor.

   b) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

9. In the navigation pane, click **Memory**.

   a) In the **Memory** page, you can view the properties of the VIOS that is using dedicated or shared memory. You can also allocate the required amount of dedicated or shared memory to the VIOS. Click **Advanced** to change the Assigned Barrier Synchronization Register (BSR) Array.

      Note: POWER8® or POWER9™processor-based servers do not support BSR.

   b) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

10. In the navigation pane, click **Physical I/O Adapters**.

    a) The **Physical I/O Adapters** page lists the physical I/O adapters that are assigned to the VIOS partition with the adapter physical location code and description. Click **Add Adapter** to open the **Add Physical I/O Adapters** page. In the **Add Physical I/O Adapters** page, select a drawer to list the available adapters or filter the adapters by their physical location. Select an adapter from the table and click **OK**. Right-click an adapter in the **Physical I/O Adapter** page and select **Remove Adapter** to remove an adapter after confirmation.

    b) Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

11. In the navigation pane, expand **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page is displayed with the **Virtual SCSI Adapters** and **Virtual Fibre Channel Adapters** tabs. By default, the **Virtual SCSI Adapters** tab is selected.

    a) **Virtual SCSI Adapters**

       i) In the **Virtual SCSI Adapters** section, click **Create Adapter**. The **Create Virtual SCSI Adapter** window is displayed.

       ii) In the **Server Adapter ID** field, enter the server adapter ID.

       Note: If you do not want to specify the server adapter ID, you can continue the procedure with the server adapter ID that is populated automatically in the **Server Adapter ID** field. The server adapter ID displayed in this field is the next available slot ID for the virtual SCSI server adapter that is being created.

       iii) From the **Remote Partition** list, select the logical partition to which the virtual SCSI adapter connects. The list displays all the logical partitions that are available in the managed system for creating the virtual SCSI adapter.

       iv) From the **Remote Adapter ID** list, select the remote adapter ID. The remote slot number of the selected logical partition is displayed in the **Remote Partition ID** field. This field is populated automatically with the next available slot ID which is based on the logical partition that is selected for creating the virtual SCSI adapter. Alternatively, you can click **Populate existing usable Remote Adapter IDs**. All the client adapters that exist in the selected logical partition, and which are not connected to any Virtual I/O Server are displayed in the **Remote Adapter ID** field.

       v) By default, both the virtual SCSI **server adapter** and the corresponding **client adapter** will be created. If you do not want to create the virtual SCSI **client adapter**, clear the **Create Remote Adapter** check box.

       vi) Click **OK** to apply the changes. Alternatively, you can click **Cancel** to reject the changes and to close the window.

b) **Virtual Fibre Channel Adapters**

   i) In the **Virtual Storage** page, select **Virtual Fibre Channel Adapters** tab.

   ii) In the **Virtual FC Adapters** section, click **Create Adapter**. The **Create Virtual Fibre Channel Adapter** window is displayed.

   iii) In the **Server Adapter ID** field, enter the server adapter ID.

   Note: If you do not want to specify the server adapter ID, you can continue the procedure with the server adapter ID that is populated automatically in the **Server Adapter ID** field. The server adapter ID displayed in this field is the next available slot ID for the virtual Fibre Channel server adapter that is being created.

   iv) From the **Remote Partition** list, select the logical partition to which the virtual Fibre Channel adapter connects. The list displays all the logical partitions that are available in the managed system for creating the virtual Fibre Channel adapter.

   v) From the **Remote Adapter ID** list, select the remote adapter ID. The remote slot number of the selected logical partition is displayed in the **Remote Partition ID** field. This field is populated automatically with the next available slot ID, which is based on the logical partition that is selected for creating the virtual Fibre Channel adapter. Alternatively, you can click **Populate existing usable Remote Adapter IDs**. All the client adapters that exist in the selected logical partition, and which are not connected to any Virtual I/O Server, are displayed in the **Remote Adapter ID** field.

   vi) By default, both the virtual Fibre Channel **server adapter** and the **client adapter** will be created. If you do not want to create the virtual Fibre Channel **client adapter**, clear the **Create Remote Adapter** check box.

   vii) Click **OK** to apply the changes. Alternatively, you can click **Cancel** to reject the changes and to close the window.

12. In the navigation pane, expand **Virtual I/O** > **Hardware Virtualized I/O**. The Single Root I/O Virtualization (**SR-IOV**) page for the selected Virtual I/O Server is displayed in the work pane.

   a) The **SR-IOV** page lists all of the SR-IOV logical ports that are connected to the VIOS. Right-click a logical port and select **Modify Port** or **Remove Port** to change or remove the selected port. Click **Add Port** to add an SR-IOV logical port to the VIOS partition. Click **Select and SR-IOV physical port** to view a list of available physical ports. After you select a physical port, a table is displayed that lists the configuration details of the physical port. You can also configure additional settings for the logical port in the advanced settings section.

   b) The **HEA** page lists all of the Logical Host Ethernet Adapters (LHEAs) connected to the VIOS. Click **Add Adapter** to assign more adapters to the VIOS partition. You can modify the adapter to use dedicated resources by selecting **Yes** in the **Dedicated** column. Click **Advanced Settings** to set the Media Access Control (MAC) address settings and virtual LAN (VLAN) ID settings.

   You can select the following values for operating system-defined MAC address.

   - **Allow all**: Allows any operating-system defined MAC address. This value is the default value that is displayed.
   - **Deny all**: Does not allow any operating-system defined MAC addresses.
   - **Allow specified**: Specifies a maximum of four operating-system defined MAC addresses that are allowed. You can add the MAC addresses to the Allowed MAC Addresses list.

   You can set the adapters to accept packets with any virtual LAN ID (VLAN ID) or to accept only packets with specific VLAN IDs.

   - **Allow all**: Allows the logical port to accept packets with any VLAN ID.
   - **Deny all**: Does not allow the logical port to accept packets with any VLAN ID.
   - **Allow specified**: Allows the logical port to accept packets with only specific VLAN IDs.

   Note:

You must use the following configuration settings when you specify the MAC address and VLAN ID for the configuration to be valid:

- If **MAC Address Settings** is set to **Allow all**, the **VLAN ID Settings** must also be set to **Allow all**. Any other value that is specified for the VLAN ID is not valid.
- If **MAC Address Settings** is set to **Deny all**, the **VLAN ID Settings** can either be set to **Deny all** or **Allow Specified**.
- If **MAC Address Settings** is set to **Allow Specified**, the **VLAN ID Settings** can either be **Deny all** or **Allow Specified**.

c) In the **HCA** page click **Launch Manage Host Channel Adapters** to open the HMC panel with a list of available Host Channel Adapters (HCA). Select an HCA to display the current partition usage for the selected HCA.

Managing Virtual I/O Server operations

You can shut down or restart a Virtual I/O Server (VIOS) by using the Hardware Management Console (HMC).

For instructions, see Shutting down a Virtual I/O Server and Restarting a Virtual I/O Server.

Accessing management operations for a VIOS

You can use the Hardware Management Console (HMC) to manage a Virtual I/O Server (VIOS).

About this task

To access the management operations for a VIOS, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, click the server name that has the VIOS.
4. In the **PowerVM** area click **Virtual I/O Servers**. The Virtual I/O Servers that are available on the system are displayed.
5. Select the VIOS and click **Actions** > **View Virtual I/O Server Properties**
6. In the work pane, select the Virtual I/O Server of your choice and select a management task from the options.

Changing the default profile of a VIOS

You can change the default profile of a Virtual I/O Server (VIOS) by using the Hardware Management Console (HMC).

About this task

To change the default profile of a VIOS by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.

4. In the **PowerVM** area, click **Virtual I/O Servers** to view all the Virtual I/O Servers on the selected system.
5. In the work pane, right-click the Virtual I/O Server of your choice and select **Profiles** > **Change Default Profile**. The **Change Default Profile** page is displayed.
6. From the **New Default Profile** list, select a new default profile.

Managing virtual networks

Learn about IBM PowerVM networking concepts and managing PowerVM virtual networks.

The IBM Power Architecture® defines a set of networking technologies with specific terminology. You can use the Hardware Management Console (HMC) to manage PowerVM virtual networks.

PowerVM networking concepts

PowerVM includes extensive and powerful networking tools and technologies, which you can use to enable more flexibility, better security, and enhanced usage of hardware resources. Some of these terms and concepts are unique to the Power Architecture.

Network connectivity in the PowerVM virtual environment is highly flexible. PowerVM virtual networking includes the following technologies:

Table 1. PowerVM network technologies

PowerVM technology	Definition
Virtual network	Enables interpartition communication without assigning a physical network adapter to each partition. If the virtual network is bridged, partitions can communicate with external networks. A virtual network is identified by its name or VLAN ID and the associated virtual switch.
Virtual Ethernet adapter	Enables a client partition to send and receive network traffic without a physical Ethernet adapter.
Virtual switch	An in-memory, hypervisor implementation of a layer-2 switch.
Virtual network bridge	A software adapter that bridges physical and virtual networks to enable communication. A network bridge can be configured for failover or load sharing.
Link aggregation device	A link aggregation (also known as Etherchannel) device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated.

Virtual networks

The managing PowerVM option includes the **Add Virtual Network** wizard that guides you through the steps to create the virtual network. A PowerVM virtual network allows connectivity between partitions on a server or, if bridged, across servers. You can create multiple virtual networks on a managed system and then connect partitions to those networks.

A virtual local area network (VLAN) allows the physical network to be logically segmented. You can connect partitions to virtual Ethernet adapters, and then connect those adapters to VLANs. Traffic on the VLANs can be routed through virtual switches.

A VLAN is a method to logically segment a physical network so that layer 2 connectivity is restricted to members that belong to the same VLAN. This separation is achieved by tagging Ethernet packets with their VLAN membership information and then restricting delivery to members of that VLAN. VLAN is described by the IEEE 802.1Q standard.

The VLAN tag information is referred to as VLAN ID (VID). Ports on a switch are configured as being members of a VLAN designated by the VID for that port. The default VID for a port is referred to as the Port VID (PVID). The VID can be added to an Ethernet packet either by a VLAN-aware host, or by the switch in the case of VLAN-unaware hosts. Therefore, ports on an Ethernet switch must be configured with information that indicates whether the host connected is VLAN-aware.

For VLAN-unaware hosts, a port is set up as untagged and the switch tags all packets that enter through that port with the Port VLAN ID (PVID). The switch also untags all packets that exit that port before delivery to the VLAN unaware host. A port that is used to connect VLAN-unaware hosts is called an *untagged port*, and it can be a member of only a single VLAN identified by its PVID. Hosts that are VLAN-aware can insert and remove their own tags and can be members of more than one VLAN. These hosts are typically attached to ports that do not remove the tags before the packets are delivered to the host. However, it inserts the PVID tag when an untagged packet enters the port. A port allows only packets that are untagged or tagged with the tag of one of the VLANs that the port belongs to. These VLAN rules are in addition to the regular media access control (MAC) address-based forwarding rules followed by a switch. Therefore, a packet with a broadcast or multicast destination MAC is also delivered to member ports that belong to the VLAN that is identified by the tags in the packet. This mechanism ensures the logical separation of the physical network that is based on membership in a VLAN.

Virtual Ethernet adapters

A virtual Ethernet adapter allows client partitions to send and receive network traffic without a dedicated physical Ethernet adapter. A virtual Ethernet adapter is created when you connect a partition to a virtual network. You can change and connect the virtual Ethernet adapters to virtual networks. TCP/IP communications over these virtual networks are routed through the server firmware at high speed.

Virtual Ethernet adapters allow logical partitions within the same system to communicate without having to use physical Ethernet adapters. Within the system, virtual Ethernet adapters are connected to an IEEE 802.1Q virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VIDs. With VIDs, virtual Ethernet adapters can share a common logical network. The system transmits packets by copying the packet directly from the memory of the sender logical partition to the receive buffers of the receiver logical partition without any intermediate buffering of the packet.

You can use virtual Ethernet adapters without using the Virtual I/O Server, but the logical partitions cannot communicate with external systems. However, in this situation, you can use another device, called a Host Ethernet Adapter (or Integrated Virtual Ethernet) to facilitate communication between logical partitions on the system and external networks.

Related links

Virtual Ethernet adapters

Virtual Ethernet

Virtual local area networks

Virtual switches

The POWER Hypervisor implements an IEEE 802.1Q virtual LAN style virtual Ethernet switch. When you add a virtual network, you can add a virtual switch. After you add a virtual switch, if necessary, you can change the name and mode of the virtual switch.

Multiple virtual switches are supported. By default, a single virtual switch that is named *ETHERNET0* is configured. You can change the name of the virtual switch and create more virtual switches with different names by using the Hardware Management Console (HMC). You can add more virtual switches to provide an extra layer of security or to increase the flexibility of a virtual Ethernet configuration.

Note: A virtual switch that is associated with a virtual network bridge can be removed only if the following conditions are true:

- All virtual network bridges to which the virtual switch is attached are deleted.
- The virtual switch is not associated with any other virtual network bridge.

Related links

Changing system configuration

Changing the virtual switch mode setting

Configuring the Virtual I/O Server for the VSN capability

Virtual network bridges

A virtual network bridge can be configured for a failover or for load sharing. If the virtual network bridge is configured for a failover, a primary Virtual I/O Server (VIOS) and a backup VIOS must be identified.

A virtual network bridge has one or more load groups. By default, a virtual network bridge has one load group. The number of load groups determines the number of virtual Ethernet adapters (trunk adapters) present on each Shared Ethernet adapter (SEA) that is a part of the virtual network bridge.

The PowerVM virtual network bridge is associated with one or more shared Ethernet adapters (SEAs) that bridge the internal network traffic to a physical network adapter. You can create or change a network bridge for the virtual networks by using the Hardware Management Console (HMC).

A virtual network that is connected through a virtual network bridge can be tagged or untagged. If you are creating a tagged network, you can choose an existing network bridge or create a network bridge for the virtual network that you want to add to the managed system. If you are creating an untagged network, you must create a new network bridge. In an untagged virtual network, PowerVM uses the virtual LAN ID to tag and route the network traffic among partitions.

A virtual network bridge can be associated with one untagged virtual network and up to 20 tagged virtual networks. A bridged virtual network is created by adding a virtual network to an existing or a new virtual network bridge. When a virtual network is added to an existing bridge, a tagged virtual network is created. When a virtual network is added to a new bridge, it can be added as an untagged network or as a tagged network.

A virtual network bridge can be configured for a failover only when two Virtual I/O Servers are present in the network. Also, each VIOS must have only one trunk adapter that is associated with a particular VLAN configuration. The priority of the trunk adapter must be unique for a specific VLAN configuration. You cannot configure a virtual network bridge with more than two trunk adapters by using the HMC Graphical User Interface (GUI) or the HMC REST API. However, if required, you can use the HMC command-line interface (CLI) and VIOS commands to create more than two trunk adapters, with the same VLAN configuration. Additionally, you can set different priorities for trunk adapters across Virtual I/O Servers. After creating trunk adapters, the HMC REST API or the HMC GUI does not support any operation on that virtual network bridge. You must use the HMC CLI and VIOS commands to delete any trunk adapters and continue the operation on that virtual network bridge by using the HMC REST API or the HMC GUI.

Link aggregation devices

A link aggregation, or Etherchannel device, is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters that are aggregated can then act as a single Ethernet device. Link aggregation helps provide more throughput over a single IP address than would be possible with a single Ethernet adapter.

For example, `ent0` and `ent1` adapters can be aggregated to the `ent3` adapter. The system considers these aggregated adapters as one adapter, and all adapters in the link aggregation device are given the same hardware address. Therefore, they are treated by remote systems as if they were one adapter.

Link aggregation can provide increased redundancy because individual links might fail. The link aggregation device can automatically fail over to another adapter in the device to maintain connectivity. For example, if the `ent0` adapter fails, the packets are automatically sent on the next available adapter, `ent1`, without disruption to existing user connections. The `ent0` adapter automatically returns to service on the link aggregation device when it recovers.

Related information

Network attributes

Viewing the virtual network configuration

On a server that is managed by the Hardware Management Console (HMC), you can view the configuration details of the PowerVM virtual networks.

Procedure

To view and change resources and the network configuration for a Virtual I/O Server (VIOS), complete the following steps:

1. In the navigation pane, click the **Resources** icon
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the **Virtual Networks** work pane, you can use the left and right arrow key buttons to switch between **Network(s)** and **Adapter(s)** views. The **Network(s)** view lists all the virtual networks that are configured on the managed system. Each table represents the properties of the Virtual Networks, Virtual Switches, Network Bridges, and Link Aggregation Devices. The **Adapter(s)** view lists all the network adapters that are connected to the logical partition. You can view the Virtual I/O Servers and the associated virtual Ethernet adapter IDs, load group, VLAN ID, and the 802.1Q VLAN ID settings for the adapter in the table.
   a) The **Virtual Networks** section lists all the virtual networks that are configured on the managed system. Virtual network is a system level attribute that helps you to create multiple virtual networks on the managed system. Right-click a virtual network in the table and select **Modify virtual network name** to change the virtual network name. Select **View connected partitions** to view the partitions that are connected to the selected virtual network. Select **Remove virtual network** to remove the virtual network from the partition after confirmation. Click **Add Virtual Network** to add a network to the partition by using the **Add Virtual Network** wizard.
   b) The **Virtual Switches** section lists all the virtual switches that are configured on the managed system. A virtual switch (VSwitch) is used to allow the virtual Ethernet adapters to route through a physical adapter to an external network. Right-click a virtual switch in the table and select **Modify virtual switch** to change the virtual switch name. Select **Remove virtual switch** to remove the virtual switch from the partition after confirmation.
   c) The **Virtual Network Bridges** section lists all the virtual network bridges that are configured on the managed system. A network bridge is used to associate one or more shared Ethernet adapters that bridge internal network traffic to a physical network adapter. Right-click virtual network bridge in the table and select **Modify virtual network bridge**, or **View virtual network bridge** to change the properties of the selected virtual network bridge, or select **Add Virtual Network to Load Group** to add a virtual network bridge to a load group.
   d) The **Link Aggregation Devices** section lists all the link aggregation devices on the VIOS. Right-click a device in the table and select **Modify** or **Remove** to change the properties of the selected device. Click **Add Device** to add a link aggregation device. Select a VIOS and mode for the device.

Results

You can view the configuration details of the virtual networks in the table that is displayed in the **Virtual Networks** tab. The configuration details for each virtual network include the following information:

- Virtual network name
- VLAN ID
- Virtual switch
- Virtual network bridge

- Load group

The Add Virtual Network wizard

You can use the **Add Virtual Network** wizard in the Hardware Management Console (HMC) to add an existing virtual network or to add a new virtual network to the server.

You can complete the following tasks by using the **Add Virtual Network** wizard:

- Create internal or bridged networks.
- Create tagged or untagged virtual networks.
- Create a virtual network on an existing or a new virtual switch.
- Create a load group or select an existing load group.

Note: When you add a virtual network, the wizard prompts you to create a network bridge to support the new virtual network. You can connect the new virtual network to an existing network bridge or create a network bridge. If you select the untagged network, you are prompted to create a new network bridge. If the physical network adapters are unavailable to create a network bridge, you cannot select an untagged network.

Adding a virtual network with an existing virtual network bridge

On a server that is managed by the Hardware Management Console (HMC), you can add a PowerVM virtual network with an existing virtual bridge by using the **Add Virtual Network** wizard.

Procedure

To add a virtual network with an existing virtual bridge on a server by using the **Add Virtual Network** wizard, complete the following steps:

1. In the navigation pane, click the **Resources** icon
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, click **Add Virtual Network**. The **Add Virtual Network** wizard opens to the **Network Name** page.
   a) Enter a name in the **Virtual network name** field.
   b) Select **Bridged Network** or **Internal Network** to specify the type of virtual network.
   c) Select **Yes** from the **IEEE 802.1Q Tagging** list to specify that the network is tagged.
   d) Enter a virtual network ID in the **VLAN ID** field. The valid range for the ID is 1 - 4094.
   e) Click **Advanced Settings** to expand the section.
   f) Select **Use an existing Virtual Switch**.
   g) Choose an existing virtual switch from the table.
   h) Select the **Add new virtual network to all Virtual I/O Servers** to add the new virtual network to all the Virtual I/O Servers.

      A client virtual Ethernet adapter is added to all the Virtual I/O Servers.
   i) Click **Next** and then continue with step 6.
6. To use an existing virtual network bridge, complete the following steps:
   a) If you want to enable failover, select **Yes** for failover from the **Network Bridge Settings** option.
   b) If you want to enable load sharing, select **Yes** for load sharing from the **Network Bridge Settings** option.
   c) Enter a network bridge PVID in the **Bridge PVID** field.

   d) Select **Jumbo Frame**, **Large Send**, and **QoS** for the **Network Bridge Settings**.
   e) Click **Next** and then continue with step 7.
7. To select the VIOS and adapter, complete the following steps:
   a) Select the Virtual I/O Server and physical adapter location as the primary Virtual I/O Server.
   b) Use the **Advanced VIOS Settings** to configure the address to ping, IP address, netmask, and the gateway details for the selected VIOS.
   c) Click **Next** and then continue with step 8.
8. To use an existing load group, complete the following steps:
   a) Select the **Use an existing Load Group**.
   b) From the table that lists the existing load group, select a load group.
   c) Click **Next** and then continue with step 9.
9. To create a load group, complete the following steps:
   a) Select the **Create a new Load Group** option.
   b) Enter a VLAN ID for the load group in the **Enter Load Group PVID** field.
   c) Enter a name for the load group in the **Load Group Name** field. A load group creates a pair of trunk adapters with the VLAN ID you enter.
   d) Click **Next** and then continue with step 10.
10. To view the summary of the virtual network that is created by using the **Add Virtual Network** wizard, complete the following steps:
   a) Click **Adapter View** or **Network View** to display the summary of the virtual network. You can use the **Adapter View** tab to change the adapter ID.
   b) Click **Finish** to exit the **Add Virtual Network** wizard.

Adding a virtual network by creating a virtual network bridge

On a server that is managed by the Hardware Management Console (HMC), you can use the **Add Virtual Network** wizard to add a PowerVM virtual network.

Procedure

To add a virtual network by creating a virtual network bridge, by using the **Add Virtual Network** wizard, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, click **Add Virtual Network**. The **Add Virtual Network** wizard opens to the **Network Name** page.
   a) Enter a name in the **Virtual network name** field.
   b) Select either **Bridged Network** or **Internal Network** depending on what type of network you want to create.
   c) Select **No** from the **IEEE 802.1Q Tagging** list to specify that the network is untagged.
   d) Enter a virtual network ID in the **VLAN ID** field. The valid range for the ID is 1 - 4094.
   e) Click **Advanced Settings** to expand the section.
   f) Select **Create a new Virtual Switch**.
   g) Enter a virtual switch name and mode for the new switch.

   h) Select **Add new virtual network to all Virtual I/O Servers** to add the new virtual network to all the Virtual I/O Servers.

      A client virtual Ethernet adapter is added to all of the Virtual I/O Servers. The VLAN ID of the virtual Ethernet adapter that is added also provides the name of the Virtual Network ID.

   i) Click **Next** and then continue with step 6.

6. To select a **Virtual Network Bridge**, complete the following steps:

   a) If you want to enable failover, select **Yes** for failover from the **Network Bridge Settings** option.

   b) If you want to enable load sharing, select **Yes** for load sharing from the **Network Bridge Settings** option.

   c) Enter a network bridge PVID in the **Bridge PVID** field.

   d) Select **Jumbo Frame**, **Large Send**, and **QoS** for the **Network Bridge Settings**

   e) Click **Next** and continue with step 7.

7. To select the VIOS and Adapters, complete the following steps:

   a) Select the Virtual I/O Server and the physical adapter location as the primary Virtual I/O Server.

   b) Use the **Advanced VIOS Settings** tab to configure the address to ping, IP address, netmask, and the gateway details for the selected VIOS.

   c) Click **Next** and continue with step 8.

8. To use an existing load group, complete the following steps:

   a) Select **Use an existing Load Group**.

   b) From the table that lists the existing load group, select a load group.

   c) Click **Next** and continue with step 9.

9. To create a load group, complete the following steps:

   a) Select the **Create a new Load Group** option.

   b) Enter a VLAN ID for the load group in the **Enter Load Group PVID** field.

   c) Enter a name for the load group in the **Load Group Name** field. A load group creates a pair of trunk adapters with the VLAN ID you enter.

   d) Click **Next** and continue with step 10.

10. To view a summary of the virtual networks, complete the following steps:

    a) Click **Adapter View** or **Network View** to display a summary of the virtual network. You can use the **Adapter View** tab to change the adapter ID.

    b) Click **Finish** to exit the **Add Virtual Network** wizard.

Changing the name of a virtual network

On a server that is managed by the Hardware Management Console (HMC), you can change the name of a PowerVM virtual network.

About this task

To change the name of a virtual network, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.

5. In the work pane, right-click the virtual network that you want to change and select **Modify Virtual Network Name**. The **Modify Virtual Network Name** page opens.
6. Change the name of the virtual network in the **Virtual network name** field.
7. Click **OK** to apply the changes.

Changing the load group of a virtual network

From the Hardware Management Console (HMC), you can change the load group of a PowerVM virtual network.

About this task

To change the load group of a virtual network, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, right-click the virtual network that you want to change and select **Modify Load Group**. The **Modify Load Groups** page is displayed.
6. Select the load sharing group that you want from the **Load Groups** table that is displayed.
7. Click **OK** to apply the changes.

Removing a virtual network

From a server that is managed by the Hardware Management Console (HMC), you can remove a PowerVM virtual network.

About this task

Important: Before you remove a virtual network, update the information of the virtual network in the list of networks if the partitions are connected. Consider the following points when you remove a virtual network:

- If the network is a tagged virtual network, remove the virtual network from the network bridge.
- If the network is either untagged or the last tagged virtual network in the bridge, remove the network bridge along with the virtual network.

To remove a virtual network, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, right-click the virtual network that you want to remove and select **Remove Virtual Network**.

Attention: A virtual network bridge associated with a virtual network can be deleted only if the following conditions are true:

- The virtual network to which the virtual network bridge is attached is deleted.
- The virtual network bridge is not associated with any other virtual network.

6. Click **OK** to remove the selected virtual network.

Changing a virtual switch

From a server that is managed by the Hardware Management Console (HMC), you can change the attributes of a PowerVM virtual switch.

About this task

To change a virtual switch, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Virtual Switches**.
6. Right-click the virtual switch that you want to change and select **Modify Virtual Switch**. Or, you can select the virtual switch and click **Action** > **Modify Virtual Switch**.
7. Change the name of the virtual switch in the **Virtual Switch Name** field.
8. Change the mode of the virtual switch to virtual Ethernet bridging (VEB) or virtual Ethernet port aggregator (VEPA).

   Note: The VEPA mode option is available on VEPA-capable hardware only.
9. Click **OK** to apply the changes.

Removing a virtual switch

From a server that is managed by the Hardware Management Console (HMC), you can remove a PowerVM virtual switch.

About this task

To remove a virtual switch, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Virtual Switches**.

6. Right-click the virtual switch that you want to remove and select **Remove Virtual Switch**. Or, you can select the virtual switch and click **Action** > **Remove Virtual Switch**.
7. Click **OK** when you are prompted to confirm the removal.

Synchronizing a virtual switch

From a server that is managed by the Hardware Management Console (HMC), you can synchronize a PowerVM virtual switch.

About this task

To synchronize a virtual switch, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Virtual Switches**.
6. Right-click the virtual switch that you want to synchronize and select **Sync Virtual Switch**. Or, you can select the virtual switch and click **Action** > **Sync Virtual Switch**.
7. Click **OK** when you are prompted to confirm the synchronization.

Changing a network bridge

From a server that is managed by the Hardware Management Console (HMC), you can change the PowerVM virtual network bridge properties.

Procedure

To change virtual network bridge properties, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Virtual Network Bridges**.
6. Right-click the virtual network bridge that you want to change and select **Modify Network Bridge**.
7. Enable or disable network failover in the **Failover** field.
8. Enable or disable load sharing in the **Load Sharing** field.
9. Change the physical adapter location for the primary Virtual I/O Server (VIOS) in the table.
10. Enable **Jumbo Frame** on the network bridge for the virtual Ethernet adapter to communicate to an external network.

    Note: Before you enable jumbo frames on a network bridge, check whether other devices in your network are also configured for jumbo frames.
11. Enable **Large Send** on the network bridge to reduce the processor usage of the VIOS.

12. Enable **QoS** on the network bridge to check the priority value of all tagged packets and arrange those packets in the corresponding queue.
13. Click **OK** to apply the changes.

Adding a link aggregation device

On a server that is managed by the Hardware Management Console (HMC), you can add a link aggregation device to the VIOS by using the **Add Link Aggregation device** wizard.

About this task

Note: Ensure that the VIOS is assigned with one or more physical Ethernet adapters and at least one link aggregation interface exists on the VIOS.

Procedure

To add a link aggregation device, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Link Aggregation Devices** and click **Add device**.
6. Select the Virtual I/O Server.
7. Set the mode as **standard**, **IEEE 802.3 AD**, or **round robin**.
8. Select the port location from the table in the **Port Physical Location** field.
9. Click **OK** to apply the changes.

Changing a link aggregation device

From a server that is managed by the Hardware Management Console (HMC), you can change the properties of a link aggregation device.

Procedure

To change the properties of a link aggregation device, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Link Aggregation Devices**.
6. Right-click the device that you want to change and select **Modify Link Aggregation Device**.
7. Change the mode of the device in the **Mode** field.
8. Change the port location in the **Port Physical Location** field. You can also select more than one port location or disable selected port locations.
9. Click **OK** to apply the changes.

Removing a link aggregation device

From a server that is managed by the Hardware Management Console (HMC), you can remove a link aggregation device.

Procedure

To remove a link aggregation device, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Networks**. The **Virtual Networks** page opens.
5. In the work pane, expand **Link Aggregation Devices**.
6. Right-click the device that you want to remove and select **Remove**.
7. Click **OK** to remove the device.

Managing virtual Network Interface Controllers

A virtual Network Interface Controller (vNIC) is a type of virtual Ethernet adapter that is configured on client partitions of Power Systems servers. Each vNIC is backed by an SR-IOV logical port that is available in a Virtual I/O Server (VIOS) partition. This type of vNIC is also called dedicated vNIC, as the backing SR-IOV logical port serves the vNIC exclusively. The key advantage of placing the SR-IOV logical port in the VIOS is that it makes the client LPAR eligible for Live Partition Mobility (LPM). Although the backing device esists remotely, through a mature PowerVM technology that is known as Logical Redirected DMA (LRDMA), the vNIC can map its transmit and receive buffers to the remote SR-IOV logical port when a one-to-one relationship exists between the vNIC and the backing logical port. After the buffers are mapped, the SR-IOV logical port directly fetches/stores packet data from/to the memory of the client partition. The LRDMA technology eliminates two data copies incurred in the traditional virtual Ethernet that is backed by Shared Ethernet Adapter, thus lowering CPU and memory consumption on the VIOS. Furthermore, because of the one-to-one relationship, the resources that are provisioned for the SR-IOV logical port are owned by the vNIC. As a result, the vNIC inherits all the capabilities that the SR-IOV adapter offers such as the QoS minimum-bandwidth assurance and the ability of setting PVID, VLAN ACL, and MAC ACL.

The vNIC configuration requires the following firmware and operating system support:

- System firmware level FW840 and HMC 840, or later.
- VIOS 2.2.4.0, or later.
- vNIC driver support from AIX and IBM i systems.

Dedicated vNICs backed up by SR-IOV logical ports

For dedicated vNICs, SR-IOV logical ports are the only ones that can be used as backing devices. To create a vNIC, you need to specify the hosting VIOS, in addition to the backing SR-IOV adapter and the physical port from which the logical port is to be allocated. You can also specify the VLAN settings and the MAC settings. For more information, see "Adding virtual NICs" on page 85. The VLAN settings and MAC settings are applied to both the vNIC and the SR-IOV logical ports. Default settings are applied if you do not specify the required parameters. When you add a vNIC in the client LPAR, the backing devices are provisioned and configured automatically by the HMC (based on your specification or defaults). Similar automation is performed for vNIC removal. This setup implies that you need to deal with only the client vNIC adapter and not be concerned with the management of backing devices, in normal cases.

Note:

- HMC supports the vNIC configuration in GUI, command line, and REST APIs.

- Most of the HMC GUI support for vNIC (vNIC add, delete, or edit) is available in the enhanced HMC mode only (not in the classic mode).
- HMC automated management of the backing devices requires RMC connection to the hosting VIOS.

LPM consideration for vNIC

During Live Partition Mobility (LPM) operations, HMC handles the creation of the vNIC server and backing devices on the target system and cleanup of devices on the source system, when LPM completes successfully. HMC has built-in capability to provide auto-mapping of backing devices and hosting Virtual I/O Servers between the source and target servers. The SR-IOV port label, the available capacity and the VF count, and the adapter and VIOS redundancy are some of the key factors that are used by the HMC for auto-mapping. Optionally, you can also specify your own mapping settings.

Viewing virtual NIC backing devices

You can use the Hardware Management Console (HMC) to view the virtual NIC backing devices.

About this task

To view the virtual NIC backing devices that are assigned to the Virtual I/O Server (VIOS) by using HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual NICs**. The **Virtual NIC Backing Devices** page opens with the virtual Network Interface Controllers (vNICs) listed in a table. The table lists all the devices on the managed system that are configured as backing devices for the virtual NICs. You can also view other information about the devices such as the name of the device, the partition that is associated with the virtual NIC, the location code of the backing device, port switch mode, port label, sub label, and the Virtual I/O Server to which each backing device is assigned.

Managing virtual storage

You can use the Hardware Management Console (HMC) to manage and monitor storage devices in a PowerVM virtual storage environment.

You can change the configuration of the virtual storage devices that are allocated to each Virtual I/O Server (VIOS) on the managed system. You can also add a VIOS to a shared storage pool cluster and manage all the shared storage pool clusters.

The virtual storage page has the adapter view and the storage view. You can toggle between these views by clicking the button in the upper-right corner of the work pane. The default view is the **Storage View**. You can use the storage view to view and manage the storage capability of the managed system.

You can view the adapter configuration of the virtual storage devices that are allocated to the Virtual I/O Servers. The **Adapter View** provides a mapping of the adapters to the physical storage device. By selecting a VIOS, you can manage the virtual storage devices that are configured to a particular partition. You can also select and view all the partitions with storage that is provisioned by the VIOS.

Moving an optical device to another partition

With the support of the Virtual I/O Server (VIOS), you can share a CD or DVD that is assigned to the VIOS among multiple AIX, IBM i, and Linux client partitions.

A shared optical device can be accessed only by one client partition at a time. If another client partition wants to use the shared optical device, you must first deallocate the shared optical device from the client partition that is accessing it.

For more information, see "Loading and unloading media files" on page 100.

Moving a virtual tape device to another partition

With the support of the Virtual I/O Server (VIOS) for the virtual tape devices, you can share the physical tape drive that is assigned to the VIOS partition among multiple AIX, IBM i, and Linux client partitions.

A shared tape device can be accessed only by one VIOS client partition at a time. If another VIOS client partition wants to use the shared tape device, you must first deallocate the shared tape device from the client partition that is accessing it.

For more information, see "Loading and unloading media files" on page 100.

Tracing the virtual storage configuration

You can track which virtual objects correspond to which physical objects. A single virtual server can have multiple virtual disks.

The virtual disks are mapped to physical disks as physical volumes or as logical volumes. The logical volumes are mapped from volume groups or storage pools.

Depending on the type of storage provisioning method you choose, you can track the following information:

- VIOS
  - Server host name
  - Physical disk location
  - Physical adapter device name
  - Physical hdisk device name
  - Cluster name (for shared storage pool backed devices only)
  - Volume group or storage pool name (for logical volume or storage pool backed devices only)
  - Logical volume or storage pool backing device name (for logical volume or storage pool backed devices only)
  - Virtual Small Computer System Interface (SCSI) adapter slot
  - Virtual SCSI adapter device name
  - Virtual target device
- VIOS client partition
  - Client host name
  - Virtual SCSI adapter slot
  - Virtual SCSI adapter device name
  - Virtual disk device name

Managing optical devices

You can use the Hardware Management Console (HMC) to view and to change optical devices.

You can add optical devices or remove optical devices to or from any partition if the partition is in either active or inactive state. If you remove an optical device from an active partition, the HMC prompts you to confirm the removal before you remove the optical device. To assign an optical device to a client partition, ensure that the client partition owns one or more virtual Small Computer System Interface (SCSI) adapters. Also, ensure that the Virtual I/O Server (VIOS) owns the corresponding virtual SCSI adapters that host the client adapter.

Managing physical optical devices

You can virtualize the physical optical devices that are assigned to the Virtual I/O Server (VIOS) by using the Hardware Management Console (HMC). The virtualized devices are shared among the client partitions of the VIOS.

Viewing physical optical devices

You can use the Hardware Management Console (HMC) to view the physical optical devices.

About this task

To view the physical optical devices that are assigned to the Virtual I/O Server (VIOS) by using HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window is displayed.
7. Click the **Optical Devices** tab to show a list of virtual optical media and physical optical devices on the managed system.
8. Select a physical optical device from the table that you want to view.
9. From the **Select Action** list of the **Physical Optical Devices** table, select **Properties** to view the properties of the selected physical optical device.

Changing the partition assignment for a physical optical device

You can use the Hardware Management Console (HMC) to change the Virtual I/O Server (VIOS) to which the optical device is assigned, or to set the optical device so it is not assigned to any other partition.

About this task

To change the partition assignment for a physical optical device by using HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.

7. Click the **Optical Devices** tab to show a list of virtual optical media and physical optical devices on the managed system.
8. Select an optical device from the **Physical Optical Devices** table for which you want to change the partition assignment.
9. From the **Select Action** list of the **Physical Optical Devices** table, select **Modify assignment** option. The **Modify Physical Optical Device Assignment** page is displayed.
10. Change the partition to which the optical device is assigned, or set the optical device so it is not assigned to any partition. Click **OK**. The list of optical devices reflects the changes that you made.

Managing virtual optical devices

You can virtualize a DVD or a CD device that is assigned to the Virtual I/O Server (VIOS) by using the Hardware Management Console (HMC). The virtualized devices are shared among the client partitions of the VIOS.

Only one client partition can access the shared optical device at a time. The advantage of a virtual optical device is that you do not have to move the parent Small Computer System Interface (SCSI) adapter between the VIOS client partitions. You cannot share optical devices if the SCSI adapter also controls the internal disk drives on which the VIOS is installed.

Note: You cannot move the virtual drive to another VIOS because client SCSI adapters cannot be created in a VIOS. If you want to virtualize the CD or DVD drive in another VIOS, the virtual device must be unconfigured and the parent SCSI adapter must be unconfigured and must be moved.

To change the virtual optical media, consider the following system requirements:

- The HMC must be Version 7 release 3.4.2, or later.
- The VIOS must be Version 2.1.1.0, or later.
- The resource monitoring and control (RMC) connection is established between the HMC and the VIOS.
- The virtual media library exists before you manage, create, or assign virtual optical devices.

Managing media libraries
A media library is a collection of virtual optical media. You can use the Hardware Management Console (HMC) to manage those libraries and assign resources to client partitions.

Viewing media libraries
You can use the Hardware Management Console (HMC) to view the media libraries.

About this task

To view the media libraries that are assigned to the Virtual I/O Server (VIOS) by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Click the **Optical Devices** tab to display a list of virtual optical media and physical optical devices on the managed system.

8. Select a media library from the **Virtual Optical Media** table that you want to view.
9. From the **Select Action** list of the **Virtual Optical Media** table, select **Properties** to view the properties of the selected media library.

Adding or removing a media library
You can use the Hardware Management Console (HMC) to add or remove media libraries to and from a selected Virtual I/O Server (VIOS).

About this task

To add or remove media libraries by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Select the options to either add or remove a media library.
8. Click **Apply** to apply the changes.

Adding or removing media files from a media library
You can use the Hardware Management Console (HMC) to add or remove media files to and from a media library that is assigned to a Virtual I/O Server (VIOS).

About this task

To add or remove media files from a media library by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Click the **Optical Devices** tab to display a list of virtual optical media and physical optical devices on the managed system.
8. Select a media library from the **Virtual Optical Media** table to add or remove media files.
9. From the **Select Action** list of the **Virtual Optical Media** table, select one of the following options:

   - **Add Media** adds an optical media file to the media library and makes it available for assignment to a partition.
   - **Delete** removes the selected media files from the media library.
10. Click **Apply** to apply the changes.

Changing the partition assignment for a media file
You can use the Hardware Management Console (HMC) to change the partition assignment for a media file by changing the virtual optical device to which a media file is assigned. You can assign a read-only media to more than one Virtual I/O Server (VIOS).

About this task

To change the partition assignment for a media file by using HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Click the **Optical Devices** tab to display a list of virtual optical media and physical optical devices on the managed system.
8. Select a media library from the **Virtual Optical Media** table for which you want to change the partition assignment for a media file.
9. From the **Select Action** list of the **Virtual Optical Media** table, select **Modify partition assignment** option.
10. Change the partition assignment as needed.
11. Click **Apply** to apply the changes.

Managing physical volumes

You can use the Hardware Management Console (HMC) to view and to change the assignment of physical volumes.

Viewing the properties of physical volumes

From a server that is managed by the Hardware Management Console (HMC), you can view the properties of the selected physical volume.

About this task

To view the properties of a physical volume by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.

3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Click **Physical Volumes** tab to display a list of physical volumes on the managed system.
8. Select the physical volume from the **Physical Volumes** table that you want to view.
9. From the **Select Action** list of the **Physical Volumes** table, select **Properties** to view the properties of the selected physical volume.

Changing physical volume assignments

From a server that is managed by the Hardware Management Console (HMC), you can change the partition to which the selected physical volume is assigned, or you can set the physical volume to ensure that it is not assigned to any other partition.

About this task

To change the physical volume assignment by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Click **Physical Volumes** tab to display a list of physical volumes on the managed system.
8. Select the physical volume from the **Physical Volumes** table that you want to change.
9. From the **Select Action** list of the **Physical Volumes** table, select **Modify partition assignment** to change the partition to which the selected physical volume is assigned, or to set the physical volume of the selected partition.

Viewing virtual SCSI adapters

You can view the properties of a virtual Small Computer Serial Interface (SCSI) adapter for each Virtual I/O Server (VIOS) that is configured on the managed system by using the Hardware Management Console (HMC). The view provides a mapping of the adapters to the physical storage device. By selecting a VIOS, you can manage the virtual storage devices that are configured to a particular partition. The virtual SCSI adapters tab displays the end to end mapping for the virtual SCSI that includes the server adapter, client adapter, and the storage that is used by the virtual SCSI adapter that is configured for a particular partition. You can also remove the client or server adapter that is configured for the particular partition.

To view the list of virtual SCSI adapters, complete the following steps:

1. In the navigation pane, click the **Resources** icon .

2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**.
5. In the **Virtual Storage** work pane, you can use the left and right arrow key buttons to switch between the **Storage(s)** and **Adapter(s)** views.
6. Click the right arrow key button to select **Adapter(s)** view.
7. Click and expand the **Virtual SCSI Adapters** section. The table lists the virtual SCSI adapters that are connected to the partition.

Viewing virtual Fibre Channel adapters

The N_Port ID Virtualization (NPIV) is an industry-standard technology that helps you to configure an NPIV capable Fibre Channel adapter with multiple, virtual worldwide port names (WWPNs). This technology is also called as virtual Fibre Channel. Similar to the virtual Small Computer System Interface (SCSI) function (VSCSI), virtual Fibre Channel is a method to securely share a physical Fibre Channel adapter among multiple Virtual I/O Servers.

Virtual SCSI server provides server-based storage virtualization. Storage resources can be aggregated and pooled on the Virtual I/O Server (VIOS). Two unique, virtual, WWPNs starting with the letter *c* are generated by the Hardware Management Console (HMC) for the virtual Fibre Channel client adapter. After the activation of the client partition, the WWPNs log in to the storage area network (SAN) similar to other WWPNs from a physical port.

From an architectural perspective, the key difference between the virtual Fibre Channel and the virtual SCSI is that the Virtual I/O Server (VIOS) does not act as a SCSI emulator to its client partitions. Instead, it acts as a direct Fibre Channel pass-through for the Fibre Channel protocol I/O traffic through the POWER Hypervisor. The client partitions are presented with full access to the physical SCSI target devices of a SAN disk or tape storage systems. The benefits of the virtual Fibre Channel are that the physical target device characteristics such as vendor or model information remains fully visible to the VIOS. Hence, you need not change the device drivers such as multi-pathing software, middleware such as copy services, or storage management applications that rely on the physical device characteristics.

Consider the following information when you use the virtual Fibre Channel:

- One virtual Fibre Channel client adapter per physical port per client partition. This strategy helps to avoid a single point of failure.
- For Fibre Channel (16GB/s or lesser) adapters, maximum of 64 active virtual Fibre Channel client adapters per physical port. The virtual adapters per physical port can reduce because of other VIOS resource constraints.
- For Fibre Channel (32GB/s) adapters, maximum of 255 virtual Fibre Channel client adapters per physical port. The virtual adapters per physical port can reduce because of other VIOS resource constraints.
- Maximum of 64 targets per virtual Fibre Channel adapter.
- 32,000 unique WWPN pairs per system. Removing a virtual Fibre Channel client adapter does not reclaim worldwide port names (WWPNs). You can manually reclaim WWPNs by using the **`mksyscfg`** command and **`chhwres`** command or by using the **`virtual_fc_adapters`** attribute.

For more information about the capabilities of the Fibre Channel adapter, see PCIe3 x8 2-port Fibre Channel (32 Gb/s).

To enable NPIV on the managed system, create the required virtual Fibre Channel adapters and connections as follows:

- You use the HMC to create virtual Fibre Channel adapters on the VIOS and associate them with virtual Fibre Channel adapters on the client partitions.

- You use the HMC to create virtual Fibre Channel adapters on each client partition and associate them with virtual Fibre Channel adapters on the VIOS. When you create a virtual Fibre Channel adapter on a client partition, the HMC generates a pair of unique WWPNs for the client virtual Fibre Channel adapter.
- You connect the virtual Fibre Channel adapters on the VIOS to the physical ports of the physical Fibre Channel adapter by running the **`vfcmap`** command on the VIOS CLI.

The HMC generates WWPNs based on the range of names available for use with the prefix in the vital product data on the managed system. You can get the 6-digit prefix when you purchase the managed system. The 6-digit prefix includes 32,000 pairs of WWPNs. When you remove a virtual Fibre Channel adapter from a client partition, the Power hypervisor deletes the WWPNs that are assigned to the virtual Fibre Channel adapter on the client partition. The HMC does not reuse the deleted WWPNs to generate WWPNs for virtual Fibre Channel adapters. If you require more WWPNs, you must obtain an activation code that includes another prefix that has another 32,000 pairs of WWPNs.

To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.

On a server that is managed by the HMC, you can dynamically add and remove virtual Fibre Channel adapters to and from the VIOS and from each client partition. You can also view information about the virtual and physical Fibre Channel adapters and the WWPNs by using VIOS commands.

For more information, see NPIV disk validation for Live Partition Migration.

Viewing virtual Fibre Channel ports for each VIOS

On a server that is managed by the Hardware Management Console (HMC), you can view the properties of the virtual Fibre Channel port that is assigned to a Virtual I/O Server (VIOS).

About this task

To view the properties of the virtual Fibre Channel ports for each VIOS, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window is displayed.
7. Click **Virtual Fibre Channel** tab to display a list of virtual Fibre Channel ports on the managed system.
8. Select the virtual Fibre Channel port from the **Virtual Fibre Channel** table that you want to view.
9. From the **Select Action** list of the Virtual Fibre Channel table, select **Properties** to view the properties of the selected virtual Fibre Channel port.

Changing to the virtual Fibre Channel adapter view

You can view the list of virtual resources per adapter that are configured for a Virtual I/O Server (VIOS).

Procedure

To view the list of virtual resources per adapter by using the Hardware Management Console (HMC), complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**.
5. In the **Virtual Storage** work pane, you can use the left and right arrow key buttons to switch between the **Storage(s)** and **Adapter(s)** views.
6. Click the right arrow key button to select **Adapter(s)** view.
7. Click and expand the **Virtual Fibre Channel Adapters** section. The table lists the virtual Fibre Channel adapters in the managed system.

Changing the virtual Fibre Channel port assignment

On a server that is managed by the Hardware Management Console (HMC), you can change the partition to which the selected virtual Fibre Channel port is assigned, or you can set the virtual Fibre Channel port to ensure that it is not assigned to any other partition.

About this task

To change the virtual Fibre Channel port assignment by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Virtual Storage Management** section to view and manage the VIOS partitions that are listed in a table.
6. Right-click the VIOS and select **Manage Virtual Storage**. The **Virtual Storage Management** window opens.
7. Click **Virtual Fibre Channel** tab to display a list of virtual Fibre Channel ports on the managed system.
8. Select the virtual Fibre Channel port from the Virtual Fibre Channel table that you want to change.
9. From the **Select Action** list of the Virtual Fibre Channel table, select **Modify virtual Fibre Channel port assignment** to change the partition to which the selected virtual Fibre Channel port is assigned, or to set the virtual Fibre Channel port to the selected partition.

Shared storage pool clusters

Shared Storage Pool (SSP) clusters are a feature in PowerVM Editions and were introduced in Virtual I/O Server (VIOS) Version 2.2.0.11 Fix Pack 11 Service Pack 1. It is a server-based storage virtualization method that provides distributed storage access to a VIOS for client partitions.

Note: For HMC to manage Shared Storage Pool clusters, the VIOS level must be at 2.2.3.3, or later.

A shared storage pool is a pool of storage area network (SAN) storage devices that can be used among Virtual I/O Servers. It is based on a cluster of Virtual I/O Servers and a distributed data object repository with a global namespace. Each VIOS that is part of a cluster represents a cluster node.

Shared storage pools provide the following benefits:

- Improve the usage of available storage.
- Simplify administration tasks.
- Simplify the aggregation of large numbers of disks among the Virtual I/O Servers.

Shared storage pools provide better usage of the available storage by using thin provisioning. The thinly provisioned device is not fully backed by physical storage if the data block is not in actual use.

Viewing the SSP cluster configuration

You can view the configuration details of Shared Storage Pool (SSP) clusters, by using the **PowerVM** > **Virtual Storage** area in the Hardware Management Console (HMC).

Procedure

To view the configuration details of shared storage pool clusters in the Virtual I/O Server (VIOS) by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Shared Storage Pool Cluster** section. The table lists the clusters that are associated with the managed system.

   Note: You can select **Show All Available Clusters** to display all clusters that are associated with the management console, not just the clusters associated with the managed system.
6. Right-click the cluster and select **View Cluster Details** to view the configuration details.
7. Click the arrows next to **Repository Disk**, **Number of cluster nodes**, **Physical Volume**, and **SSP Volume** to view more details.
8. Click **Close** to exit.

What to do next

Note: You can also view the configuration details of shared storage pool clusters in the VIOS by using the **All Shared Storage Pool Clusters** menu in the HMC. For instructions, see "Viewing the SSP cluster configuration by using the All Shared Storage Pool Clusters menu" on page 49.

Changing SSP clusters

You can change a Shared Storage Pool (SSP) cluster by using the Hardware Management Console (HMC).

Adding or removing a VIOS to an SSP cluster

You can add or remove a Virtual I/O Server (VIOS) to a Shared Storage Pool (SSP) cluster by using the **PowerVM** > **Virtual Storage** area in the Hardware Management Console (HMC).

About this task

By adding or removing a Virtual I/O Server (VIOS) to and from the shared storage pool cluster, you can extend the shared storage pool cluster. Shared storage pools extend storage virtualization to multiple Virtual I/O Servers on multiple IBM Power system servers.

Note: If the VIOS is not managed by this HMC, it cannot be removed because it will be disabled.

Procedure

To add or remove a VIOS, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Virtual Storage**. The **Virtual Storage** page opens.
5. In the **Virtual Storage** work pane, click and expand the **Shared Storage Pool Cluster** section. The table lists the clusters that are associated with the managed system.
6. To add a VIOS to the shared storage pool cluster that is a part of the managed system, complete the following steps:

   a) In the work pane, right-click a shared storage pool cluster from the table and select **Add/Remove Node**. The **Add Nodes/Remove Nodes** page displays the table with the list of Virtual I/O Servers.

   b) Select all the Virtual I/O Servers to be added to the shared storage pool cluster.

   c) Click **OK**.

7. To remove a VIOS from the shared storage pool cluster that is not a part of the managed system, complete the following steps:

   a) In the work pane, right-click a shared storage pool cluster from the table and select **Add/Remove Node**. The **Add Nodes/Remove Nodes** page opens.

   b) Clear the check box available near the Virtual I/O Servers to be removed from the shared storage pool cluster.

      Note: You cannot remove the VIOS nodes that are not managed by this HMC as they are disabled.

   c) Click **OK**.

What to do next

Note: You can also can add or remove a VIOS to a shared storage pool cluster, by using the **All Shared Storage Pool Clusters** menu in the HMC. For instructions, see "Adding nodes by using the All Shared Storage Pool Clusters menu" on page 51 and "Removing a node by using the All Shared Storage Pool Clusters menu" on page 53.

Managing shared processor pools

A shared processor pool is a PowerVM technology that you can use to control the amount of processor capacity that partitions can use from the available physical processors in the system.

Multiple shared processor pools is a capability that is supported on POWER6® technology, or later. This capability isolates work loads in a shared processor pool and prevents the work load from exceeding an upper limit. This capability is also useful for software license management, where subcapacity licensing is involved.

Up to 64 shared processor pools can be defined on IBM Power Systems servers that support multiple shared processor pools. A default shared processor pool is automatically defined in the managed system.

Each shared processor pool has a maximum processing units value that is associated with it. The maximum processing units define the upper boundary of the processor capacity that can be used by the set of partitions in the shared processor pool.

The system administrator can optionally allocate a number of reserved processing units to a shared processor pool. The reserved processing units represent the available processor capacity with the processor capacity entitlements of the individual partitions in the shared processor pool. The default value for the reserved processing units is `zero`.

By using the Hardware Management Console (HMC), you can complete the following tasks:

- Allocate a specific amount of the processing capacity from the shared processor pool to each partition that uses the shared processors.
- Configure the shared processor pools with a maximum processing unit value and a reserved processing unit value.
- View information about your shared processor pool and change the properties of that pool.

Note: The default shared processor pool is pre-configured. Hence, you cannot change the properties of the default shared processor pool. The maximum number of processors available to the default shared processor pool is the total number of active, licensed processors on the managed system minus the number of processors that are assigned to dedicated processor partitions that are set to not share their dedicated processors.

Changing a shared processor pool

You can view and change the shared processor pool configuration by using the Hardware Management Console (HMC).

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Shared Processor Pool**. The **Shared Processor Pool** page opens.
5. From the table, select the shared processor pool that you want to change.
6. From the **Select Action** list, select **Modify**.
7. Select one of the following options to change the properties of the selected shared processor pool:
    - **Pool Name** to change the name of the shared processor pool.
    - **Pool ID** to change the ID of the shared processor pool.
    - **Resource Processing Units** to change the value of the reserved processing unit. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped partitions within the shared processor pool.
    - **Maximum Processing Units** to change the maximum value of the processing unit. The maximum processing unit value limits the total number of processing units that can be used by the partitions in the shared processor pool.

What to do next

After this task is complete, assign partitions to the configured shared processor pools. You can assign a partition to a shared processor pool while creating the partition, or you can reassign existing partitions from their current shared processor pools to the shared processor pools that you configured.

When you no longer want to use a shared processor pool, you can unconfigure the shared processor pool by using this task to set the maximum number of processing units and reserved number of processing units to `0`. Before you can unconfigure a shared processor pool, you must reassign all partitions, which use the shared processor pool, to other shared processor pools.

Managing shared memory pools

You can manage the shared memory pool that is configured on a server by using the Hardware Management Console (HMC).

By using the HMC, you can complete the following management tasks on shared memory pools:

- Dynamically increase or decrease the size of the shared memory pool.
- Allocate a paging VIOS to the shared memory pool.
- Allocate a paging space device to the shared memory pool.
- Enable or disable the active memory de-duplication function.
- Delete a shared memory pool.

Important: You cannot delete a shared memory pool when shared memory partitions are configured to use the shared memory pool. The partitions must be removed or changed to dedicated memory partitions before you delete the shared memory pool.

If you want to increase the shared memory pool beyond the maximum pool size, first increase the maximum pool size to a value that is greater than or equal to the required new pool size. The maximum pool size can be increased dynamically.

Active memory de-duplication is a feature of the PowerVM Active Memory Sharing technology in which the memory pages with identical contents are de-duplicated in physical memory. Active memory de-duplication feature aggregates the same data in one memory position, and frees other duplicate memory blocks, thus optimizing memory use.

After you enable the Active Memory De-duplication option, all the partitions that are part of the shared memory pool use Active Memory De-duplication.

Changing a shared memory pool

You can view and change the shared memory pool configuration by using the Hardware Management Console (HMC).

Procedure

To change a shared memory pool, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Shared Memory Pool**. The **Create Shared Memory Pool** wizard opens to the **Welcome** page. If the Shared Memory Pool already exists, the **Modify Shared Memory Pool** wizard opens.
5. Click **Next**.
6. In the **General** page, you can view and change the shared memory pool size. Click **Next**.
7. In the **Paging VIOS** page, you can associate one or more paging VIOS partition to the shared memory pool. Click **Next**.
8. In the **Paging Space Device(s)** page, the table lists the paging space devices that are currently assigned to the shared memory pool. Choose one of the following steps:
   a) To allocate more devices to the memory pool, click **Select Devices**.
   b) To remove a device from the memory pool, click **Remove**.
9. Click **Next**. The **Summary** page displays the size of the shared memory pool, the maximum size of the pool, the paging VIOS assigned to the pool, and the paging space devices that are assigned to the pool.

10. Click **Finish** to apply the changes to the shared memory pool.

Managing reserved storage device pools

You can manage the reserved storage device pool that is configured on a server by using the Hardware Management Console (HMC).

Before you begin

The reserved storage pool has storage devices that are assigned to save data for partitions that are suspended, or for active partitions that are configured with shared memory. The required storage device space is approximately 110% of the configured maximum memory size of the partition.

A reserved storage device pool contains reserved storage devices, also known as paging space devices. These devices are similar to shared memory pools with memory size zero. To suspend a partition, a storage device must have a paging space.

One Virtual I/O Server (VIOS) must be associated as the paging service partition to the reserved storage device pool. Additionally, you can associate a second VIOS with the reserved storage device pool to provide a redundant path, and to provide higher availability for the paging space devices.

During a suspend operation, an HMC assigns a storage device from reserved storage device pool. It automatically selects an unused and suitable device from this pool to store partition suspend data. The reserved storage device must be available in the reserved storage device pool while suspending a partition.

Note: You must not suspend a partition when the **`alt_disk_install`** command is running in the VIOS on which the storage is provisioned for the client.

You can complete the following management tasks on the reserved storage device pool interface:

- Add a VIOS to the reserved storage device pool.
- Remove a VIOS from the reserved storage device pool.
- Add reserved storage devices to the reserved storage device pool.
- Remove reserved storage devices from the reserved storage device pool.

Important: You cannot delete a reserved storage device pool when partitions are configured to use the pool. The partitions must be removed or their configuration must be changed before you delete the reserved storage device pool.

When a shared memory pool is created, a reserved storage device pool is also created. When a shared memory pool is deleted, a reserved storage device pool is not automatically deleted.

A reserved storage device pool is created when a shared memory pool is created. You must create the reserved storage device pool to use the Partition Suspend and Resume capability where a shared memory pool is not configured.

About this task

To change or remove a reserved storage device pool, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Reserved Storage Pool**. The **Reserved Storage Pool Management** page opens. Choose one of the following steps:

   - Select one or more Virtual I/O Servers to assign to the reserved storage device pool.
   - Select reserved storage devices from the table and click **Select Device(s)** to assign a device.
   - Select reserved storage devices from the table and click **Remove** to remove the reserved storage device pool from the VIOS.

5. Click **Apply** to apply the changes.

Managing SR-IOV, HEA and HCA adapters

You can manage Single Root I/O Virtualization (SR-IOV), Host Ethernet Adapter (HEA), and Host Channel Adapter (HCA) settings on a server by using the Hardware Management Console (HMC).

Managing SR-IOV adapters

Single Root I/O Virtualization (SR-IOV) is an I/O virtualization technology that is used for the virtualization of I/O resources for individual servers. It logically divides a physical adapter port into multiple logical ports. This technology improves the scalability, flexibility, throughput, and latency performance of networking operations. SR-IOV is supported on certain combinations of Power Systems servers and adapters.

If an adapter supports SR-IOV, the SR-IOV tab is displayed. SR-IOV is an extension to the Peripheral Component Interconnect (PCI) Express specification to facilitate multiple partitions that are running simultaneously within a single system to share a PCI Express device. An SR-IOV capable adapter can be assigned to a partition to run in dedicated mode. Or, it can be owned by a hypervisor when the SR-IOV adapter is switched to shared mode. When an adapter is assigned to the hypervisor and is operating in shared mode, the adapter can be shared by multiple partitions at the same time.

Modifying SR-IOV adapters

You can change single root I/O virtualization (SR-IOV) adapter settings on a server by using the Hardware Management Console (HMC).

About this task

To change the SR-IOV adapter settings by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Hardware Virtualized I/O**.
5. In the **SR-IOV** tab, select an SR-IOV adapter from the **SR-IOV adapter** list. The properties of the selected SR-IOV adapter such as the mode, owner, configured logical ports, maximum logical ports are displayed.
6. Click **Modify SR-IOV**. The **Modify SR-IOV adapter** page opens with the configuration details of the selected SR-IOV adapter.
7. Change the mode by selecting **Dedicated mode** or **Shared mode** from the mode options.
8. If you choose **Dedicated mode**, remove all logical ports before you switch the SR-IOV adapter to dedicated mode.
9. Click **OK** to save the changes to the SR-IOV adapter settings.

Updating the SR-IOV adapter firmware

I/O adapters that are configured to run in Single Root I/O Virtualization (SR-IOV) shared-mode are managed by adapter driver firmware and adapter firmware. Both adapter driver firmware and adapter firmware for the SR-IOV adapter are downloaded with the system firmware updates. If you are updating the system firmware concurrently, the adapter driver firmware and the SR-IOV adapter firmware are not automatically activated for the adapters that are running in SR-IOV shared-mode to prevent any unexpected temporary outage of adapters that are running in SR-IOV shared-mode.

Two types of firmware are required to support adapters that are running in SR-IOV shared-mode. One type is the adapter driver firmware, which is used for configuring and managing the adapter. The second type is the I/O adapter firmware, which enables the adapter to interface with the adapter driver firmware. The following options can be used to activate the adapter driver firmware and the adapter firmware for adapters that are running in SR-IOV shared-mode:

- A system boot or reboot activates all the adapters that are in SR-IOV shared-mode to the new firmware level.
- When the adapter is enabled to run in SR-IOV shared-mode, the adapter driver firmware and the adapter firmware are activated to the firmware level that is available with the system firmware. Also, the activation is automatically performed to the SR-IOV adapters during the maintenance operation of an adapter that is in SR-IOV shared-mode. For example, when the SR-IOV adapter is stopped or replaced during a maintenance operation.
- A selective manual firmware activation of a SR-IOV adapter that is in SR-IOV shared-mode can be performed by using the Hardware Management Console (HMC) graphical user interface or the HMC command-line.

Notes:

1. You cannot use this procedure to update the firmware for adapters that can run in SR-IOV shared-mode, but are not running in the SR-IOV shared-mode.
2. The firmware for the adapter that is capable of running in the SR-IOV shared-mode, but currently is running in the dedicated mode and is assigned to a logical partition, can only be updated concurrently either by using the operating system (OS) that owns the adapter or by using the managing HMC (if the OS is AIX operating system or a VIOS, and when the Resource Monitoring and Control (RMC) is in a running state).

The firmware update process for the SR-IOV adapters is similar to the HMC update process of the other system firmware. When you update the system firmware, the system firmware update might also contain adapter driver firmware updates for the SR-IOV adapters, adapter firmware updates, or both. The firmware for the adapters that are configured to run in SR-IOV shared-mode is not activated automatically while they are running because of a temporary I/O outage that occurs when the firmware is activated. By not automatically activating the firmware immediately, you can schedule the most convenient time for this outage. The outage lasts approximately 1 minute for each adapter that is activated when you activate only the adapter driver firmware, and approximately 5 minutes for each adapter that is activated when you activate both the adapter driver firmware and the adapter firmware. The best practice is to activate both the adapter driver firmware and the adapter firmware simultaneous. You cannot activate only the adapter firmware. To activate the SR-IOV firmware on an adapter that is running in SR-IOV shared-mode, the managed system with the SR-IOV adapter must be powered on and in either the *Standby* state or the *Operating* state.

Updating the SR-IOV adapter firmware by using the graphical user interface
You can update the firmware for your adapters that are running in Single Root I/O Virtualization (SR-IOV) mode by using the graphical user interface when you are running system firmware level FW830, or later.

About this task

To update the firmware, complete the following steps, depending on the interface that you are using:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.

3. Select the server that is running the adapters that you want to update.
4. Click **Actions** > **View all actions** > **Updates** > **SR-IOV Firmware Update**.

   The **SR-IOV Firmware Update** panel is displayed.
5. Select one or more adapters that you would like to update.

   Use the **Update available** column to determine whether updates are available for an adapter. A value of **Yes** indicates that updates are available.

   Note: A temporary I/O outage occurs for each SR-IOV adapter while it is updated. The outage lasts approximately 1 minute for each adapter that is updated when you update only the adapter driver firmware, and approximately 5 minutes for each adapter that is updated when you update both the adapter driver firmware and the adapter firmware.
6. Right-click any of the selected adapters and click **Start firmware update**, then either **Update SR-IOV adapter driver firmware** or **Update SR-IOV adapter driver firmware and adapter firmware**

   If you selected multiple adapters, the process will serially update them. By clicking **Update SR-IOV adapter driver firmware and adapter firmware**, the brief outage is longer than when **Update SR-IOV adapter driver firmware**, but it installs all of the required updates at one time. You cannot install only the adapter firmware updates.

   The Status column updates according to the status of the update. The status is one of the following values:

   Pending Adapter Driver
   There is an adapter driver firmware update that is ready for installation.

   Pending Adapter Driver and Adapter
   There are both adapter driver firmware and adapter firmware updates that are available.

   Updating
   The firmware updates for the adapter are in progress.

   Update successful
   All updates were completed successfully.

   Update failed
   At least one of the updates for the specified adapter did not complete successfully.
7. Click **OK** to exit the Update SR-IOV firmware table when all of the adapters are updated, or click **Cancel** to stop any pending updates and leave the Update SR-IOV firmware table.

Updating the SR-IOV adapter firmware by using the command line (system level FW830 and later)
You can update the firmware for your adapters that are running in Single Root I/O Virtualization (SR-IOV) mode by using the command line. Select the procedure that applies, based on the version of your system firmware.

About this task

You can activate the available SR-IOV firmware updates by using the Hardware Management Console (HMC) command line. To activate firmware updates when you are running system firmware level FW830, or later, complete the following steps:

Procedure

1. To identify which SR-IOV adapters have available updates, enter the following command:

   ```
   lslic -t sriov -m system_name
   ```

   Where *machine_type_model* is the identifier of the system.

   The following information is displayed in comma-separated value format for each of the adapters that is running in SR-IOV mode:

   ```
   slot=SR-IOV-adapter-physical-location-code,active_adapter_driver_level=
   "current-adapter-driver-firmware-level",active_adapter_level="current-adapter-firmware-
   level",
   ```

```
update_available=0 (false)|1 (true),update_description="description",
install_separate=0 (false)|1 (true)
```

If the `update_available` value is 1, then updates are available for that adapter.

If updates are available, you can update the adapter driver firmware and the adapter firmware, or only the adapter driver firmware. To update the adapter driver firmware only, the adapter must support this operation, which is indicated if the `install_separate` value is 1. You can also update all of the adapters that require updates sequentially with a single command.

Note: A temporary I/O outage occurs for each SR-IOV adapter during its update. The outage lasts approximately 1 minute for each adapter that is updated when you update only the adapter driver firmware, and approximately 5 minutes for each adapter that is updated when you update both the adapter driver firmware and the adapter firmware.

2. Choose one of the following options that corresponds to the firmware that you want to update:

   - To update the adapter driver firmware and the adapter firmware for an SR-IOV adapter, enter one of the following commands. Updating the adapter driver firmware and the adapter firmware results in an I/O outage of up to 5 minutes for each adapter that is being updated.

     – This command updates the adapter driver firmware and adapter firmware for the adapter that is specified by the *-s* parameter.

     ```
     updlic -o f -t sriov -m system_name --subtype adapterdriver,adapter -s adapter_id
     ```

     – This command updates the adapter driver firmware and adapter firmware for the adapters that are specified by the *-s* parameter. You can specify multiple adapters by separating them with commas.

     ```
     updlic -o f -t sriov -m system_name --subtype adapterdriver,
     adapter -s adapter_id1,adapter_id2,...
     ```

   - To update only the adapter driver firmware for the selected SR-IOV adapter, enter the following command. Updating only the adapter driver firmware results in an I/O outage of up to 1 minute for each adapter during the update.

     – This command updates only the adapter driver firmware for the adapter that is specified by the *-s* parameter. You can specify more than one *adapter* by separating them with commas.

     ```
     updlic -o f -t sriov -m system_name --subtype adapterdriver -s adapter_id
     ```

3. To verify that the updates completed successfully, run the following command:

   ```
   lslic -t sriov -m system_name
   ```

   The output of the command displays the updated information about the SR-IOV adapters. Depending on which firmware you updated, the adapters with the updated firmware satisfy either the criteria of no available updates or the criterion of having only available adapter firmware updates. These criteria are shown in step “1” on page 42.

Updating the SR-IOV adapter firmware by using the command line (system firmware level earlier than FW830)

About this task

You can activate the available SR-IOV firmware updates by using the HMC command line. To activate firmware updates for system firmware level earlier than FW830, complete the following steps:

Procedure

1. To identify which SR-IOV adapters have available updates, enter the following command:

   ```
   startdump -m system_name -t resource -r "sriovdebug -fwinfo"
   ```

The output is sent to a dump file in the `/dump` directory that is titled `RSCDUMP.<serial_number>.<dump_id>.<timestamp>`. The contents of the file contains a section of information for each adapter that is running in SR-IOV mode. The section for each adapter is identified by its **Slot location code**. Use the following list to determine the state of the updates for each adapter that is listed.

- No updates are available for an adapter when the following conditions are met:
  - There is text at the end of the command output for that adapter that states there are no adapter driver firmware updates for the adapter in the specified location.
  - The version number that is displayed in the `Current Version running` output for that adapter is the same as the version number that is displayed in the `Adjunct Firmware image` output for that adapter.
- Adapter driver firmware updates are available for an adapter when the text at the end of the command output for the adapter states that there are adapter driver firmware updates for the adapter in the specified location.
- Adapter firmware updates are available for an adapter when the value of the `Current version running` for that adapter is not the same as the value of the `Adjunct Firmware image` for that adapter.

If updates are available, you can update the adapter driver firmware and the adapter firmware or only the adapter driver firmware. You can also update all of the adapters at the same time, or specify a single adapter to update.

Note: A temporary I/O outage occurs for each SR-IOV adapter during its update. The outage lasts approximately 1 minute for each adapter that is updated when you update only the adapter driver firmware, and approximately 5 minutes for each adapter that is updated when you update both the adapter driver firmware and the adapter firmware.

2. Choose one of the following options that corresponds to the firmware that you want to update:

   - To update the adapter driver firmware and the adapter firmware for an SR-IOV adapter, enter one of the following commands. Updating the adapter driver firmware and the adapter firmware results in an I/O outage of up to 5 minutes for each adapter that is being updated. Each adapter is updated sequentially, so that the total update time for updating all the adapters is up to 5 minutes per adapter, where each adapter is configured in SR-IOV shared mode.
     - This command updates the adapter driver firmware and adapter firmware for all of the adapters.

       ```
       startdump -m system_name -t resource -r "sriov all updateadapter"
       ```

     - This command updates the adapter driver firmware and adapter firmware only for the adapter that is specified by the *slot_location_code* parameter.

       ```
       startdump -m system_name -t resource -r "sriov slot_location_code updateadapter"
       ```

   - To update only the adapter driver firmware for the selected SR-IOV adapter or for all of your SR-IOV adapters, enter one of the following commands. Updating only the adapter driver firmware results in an I/O outage of up to 1 minute for each adapter during the update. Each adapter is updated sequentially, so that the total update time for updating all the adapters is up to 1 minute per adapter, where each adapter is configured in SR-IOV shared mode.
     - This command updates only the adapter driver firmware for the adapter that is specified by the *slot_location_code* parameter.

       ```
       startdump -m system_name -t resource -r "sriov slot_location_code update"
       ```

     - This command updates only the adapter driver firmware for all of the adapters.

       ```
       startdump -m system_name -t resource -r "sriov all update"
       ```

3. To verify that the updates completed successfully, run the following command:

```
startdump -m system_name -t resource -r "sriovdebug -fwinfo"
```

The output is sent to a dump file in the /dump directory that is titled RSCDUMP.*<serial_number>*.*<dump_id>*.*<timestamp>*. The contents of the file contains a section of information for each adapter that is running in SR-IOV mode. The section for each adapter is identified by its **Slot location code**. The output of the command displays the updated information about the SR-IOV adapters. Depending on which firmware you updated, the adapters with the updated firmware satisfy either the criteria of no available updates or the criterion of having only available adapter firmware updates. These criteria are shown in step "1" on page 43.

Viewing SR-IOV logical port settings

You can view the single root I/O virtualization (SR-IOV) logical port settings on a server by using the Hardware Management Console (HMC).

About this task

To see the SR-IOV logical port adapter settings by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Hardware Virtualized I/O**.
5. In the **SR-IOV** tab, select an SR-IOV adapter from the **SR-IOV adapter** list.
6. Select an SR-IOV adapter from the **SR-IOV adapter** list.
7. Select **Logical Ports** from the View options. A list of configured SR-IOV logical port adapter settings are displayed.
8. Right-click a logical port and select **View Logical Port**. The **View SR-IOV Logical Port** page opens. You can view all properties of the selected SR-IOV logical port.

Modifying SR-IOV physical port settings

You can change the single root I/O virtualization (SR-IOV) physical ports settings on a server by using the Hardware Management Console (HMC).

About this task

To change the settings of an SR-IOV physical port by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Hardware Virtualized I/O**.
5. In the **SR-IOV** tab, select an SR-IOV adapter from the **SR-IOV adapter** list. A list of SR-IOV physical ports that are configured for the selected SR-IOV adapter is displayed.
6. Right-click an SR-IOV physical port that you want to change and select **Modify Physical Port**. The **Modify SR-IOV Physical Port** page opens.

7. Change the label from the **Label** field.
8. Change the sublabel from the **Sub-Label** field.
9. Change the configured speed settings from the **Configured Speed** list.
10. Select **Advanced settings**.
11. Change the MTU size settings from the **MTU Size** list.
12. Change the port switch mode settings from the **Port Switch Mode** list.
13. Change the flow control settings from the **Flow Control** list.
14. Change the maximum number of logical ports that are supported from the **Maximum** field.
15. Click **OK** to save your changes to the SR-IOV physical port settings.

Host Ethernet Adapters (HEAs)

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Note: HEA is not supported on POWER9 processor-based server.

Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA without having to go through an Ethernet bridge on another logical partition.

To connect a logical partition to an HEA, you must create a logical Host Ethernet Adapter (LHEA) for the logical partition. A *logical Host Ethernet Adapter (LHEA)* is a representation of a physical HEA on a logical partition. An LHEA appears to the operating system as if it were a physical Ethernet adapter, just as a virtual Ethernet adapter appears as if it were a physical Ethernet adapter. When you create an LHEA for a logical partition, you specify the resources that the logical partition can use on the actual physical HEA. Each logical partition can have one LHEA for each physical HEA on the managed system. Each LHEA can have one or more logical ports, and each logical port can connect to a physical port on the HEA.

After you create an LHEA for a logical partition, a network device is created in the logical partition. This network device is named `ent`*X* on AIX logical partitions, `CMN`*XX* on IBM i logical partitions, and `eth`*X* on Linux logical partitions, where *X* represents sequentially assigned numbers. The user can then set up TCP/IP configuration like a physical Ethernet device to communicate with other logical partitions.

You can configure a partition so that it is the only logical partition that can access a physical port of an HEA by specifying *dedicated mode* for an LHEA that is assigned to the logical partition. When an LHEA is in dedicated mode, no other logical partitions can access the logical ports of the physical port that is associated with the LHEA that is in dedicated mode. You might want to configure a logical partition to dedicated mode in the following situations:

If you want to connect more than 16 logical partitions to each other and to an external network through a physical port on an HEA, you can create a logical port on a Virtual I/O Server and configure an Ethernet bridge between the logical port and a virtual Ethernet adapter on a virtual LAN. This allows all logical partitions with virtual Ethernet adapters on the virtual LAN to communicate with the physical port through the Ethernet bridge. If you configure an Ethernet bridge between a logical port and a virtual Ethernet adapter, the physical port that is connected to the logical port must have the following properties:

- The physical port must be configured so that the Virtual I/O Server is the dedicated mode logical partition for the physical port.
- The physical port can have only one logical port.

A logical port can communicate with all other logical ports that are connected to the same physical port on the HEA. The physical port and its associated logical ports form a logical Ethernet network. Broadcast and multicast packets are distributed on this logical network as though it was a physical Ethernet network. You can connect up to 16 logical ports to a physical port using this logical network. By extension, you can connect up to 16 logical partitions to each other and to an external network through this logical network. The actual number of logical ports that you can connect to a physical port depends upon the Multi-Core Scaling value of the physical port group. It also depends on the number of logical ports that have been created for other physical ports within the physical port group. By default, the Multi-Core Scaling value of each physical port group is set to 4, which allows four logical ports to be connected to the physical ports in the physical port group. To allow up to 16 logical ports to be connected to the physical ports in the physical port group, you must change the Multi-Core Scaling value of the physical port group to 1 and restart the managed system.

You can set each logical port to restrict or allow packets that are tagged for specific VLANs. You can set a logical port to accept packets with any VLAN ID, or you can set a logical port to accept only the VLAN IDs that you specify. You can specify up to 20 individual VLAN IDs for each logical port.

The physical ports on an HEA are always configured on the managed system level. If you use an HMC to manage a system, you must use the HMC to configure the physical ports on any HEAs belonging to the managed system. Also, the physical port configuration applies to all logical partitions that use the physical port. (Some properties might require setup in the operating system as well. For example, the maximum packet size for a physical port on the HEA must be set on the managed system level using the HMC. However, you must also set the maximum packet size for each logical port within the operating system.) By contrast, if a system is unpartitioned and is not managed by an HMC, you can configure the physical ports on an HEA within the operating system as if the physical ports were ports on a regular physical Ethernet adapter.

HEA hardware does not support half-duplex mode.

Managing Host Ethernet Adapters (HEAs)

You can create or change a Host Ethernet Adapter (HEA) by using the Hardware Management Console (HMC).

About this task

You can complete the following management tasks on an HEA:

- Changing an HEA adapter
- Changing an HEA port
- Viewing the partitions that are associated with an HEA port

To manage HEA tasks, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the **PowerVM** area, click **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. In the work pane, click the **HEA** tab.
6. To change an HEA adapter, complete the following steps:
   a) Select an HEA adapter from the list to display the port configuration.
   b) Click **Modify HEA Adapter**. The **Modify HEA Adapter** page opens. You can change the properties of the selected adapter, such as the Multi-Core Scaling (MCS) value for the port group. You can also view details about the port group ID, maximum logical ports, and configured logical ports.
   c) From the **HEA Port Groups** table, select a **Port Group MCS** from the list to change the MCS value.
   d) Click **OK**.
7. To change an HEA port, complete the following steps:
   a) Select an HEA adapter from the list to display the port configuration.

    b) Right click and select **Modify Port**. The **Modify HEA Port** page opens.
    c) The properties of the selected adapter port are listed. You can change the port speed, the actual maximum packet size that can be received by each physical port, and the duplex level for each physical port.
    d) Click **OK** to apply the changes.

8. To view partitions that are associated with an HEA port, complete the following steps:
    a) Select an HEA adapter from the list to display the port configuration.
    b) Right click and select **View Partitions**. The **View HEA Port Partition Assignments** page opens, which displays the partitions table that lists the partitions that are assigned to the physical port.
    c) Click **OK**.

Managing Host Channel Adapters (HCAs)

Host Channel Adapters (HCAs) provide port connections from a managed system to other devices. You can connect the port to another HCA, a target device, or a switch that redirects the incoming data from one port to a device that is attached to another port.

Before you begin

You can view a list of the HCAs on a server that is managed by the Hardware Management Console (HMC). You can select an HCA from the list to display the current partition usage for the HCA.

About this task

To view the current partition usage, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Properties**. You can view and change the properties of the system that are listed under the **PowerVM** area.
4. In the navigation pane, click **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page is displayed.
5. In the work pane, click the **HCA** tab.
6. Click **Launch Manage Host Channel Adapters**. The HMC pane opens with the list of HCA in a table.
7. From the table, select an HCA to display the current partition usage for the selected HCA.
8. Click **OK**.

Managing SSP clusters by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can use the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC) to perform management tasks for the Shared Storage Pool (SSP) clusters in the Virtual I/O Server (VIOS).

To view the configuration details of shared storage pool clusters in the Virtual I/O Server (VIOS), by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon.

2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), and the information about tiers and nodes assigned to them.
3. In the upper right of the window, click **Display Gallery View** or **Display Table View** to toggle between the table view and the gallery view.

You can manage the clusters that are listed or add additional clusters to your managed system. Select a cluster in the table to view the manage tasks and to remove the cluster from the table.

Viewing the SSP cluster configuration by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can view the configuration details of Shared Storage Pool (SSP) clusters by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To view the configuration details of shared storage pool clusters in the Virtual I/O Server (VIOS), by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), and the information about tiers and nodes assigned to them.
3. Select a Shared Storage Pool cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.

   You can view the details of the tiers, the repository disk, and the nodes assigned to the cluster. From the cluster configuration page, you can replace the assigned repository disk, add or remove nodes, and perform the following actions on the assigned tiers:

   - Add a tier
   - Remove a tier
   - Remove the default tier
   - Rename a tier
   - Set a tier as default
   - Add capacity to the tier
   - Remove capacity from the tier
   - Enable mirroring
   - Disable mirroring
   - Modify threshold percentage
   - Restrict or unrestrict system tier
4. Click **Close**.

Adding an SSP cluster by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can add Shared Storage Pool (SSP) clusters by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To add shared storage pool clusters to a Virtual I/O Server (VIOS), by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC).
3. Click **Add Shared Storage Pool Cluster**. The **Add Shared Storage Pool Cluster Wizard** opens.
4. Click **General Settings** tab.
   a) Enter a cluster name in the **Cluster Name** field.
   b) Enter a shared storage pool name in the **Shared Storage Pool** field.
   c) Under **Tier Capability**, select **Single tier capable** or **Multiple tier capable** to specify if you want the cluster to be single-tier or multitier capable. Multiple tier support provides a selection of resources, including Virtual I/O Servers that provide this feature. You can create the cluster and the system tier by using this wizard.
   d) Enter a tier name in the **System Tier Name** field.
   e) Enter the free space threshold percentage in the **Freespace Threshold %** field.
   f) Enter the overcommit threshold percentage in the **Overcommit Threshold %** field.
5. Click **Next** or click the **Nodes** tab.
   a) Select a node from the **Virtual I/O Server Cluster Nodes** table.
6. Click **Next** or click the **Repository Disk** tab.
   a) Select a disk from the **Cluster Repository Disks** table.
7. Click **Next** or click the **System Tier** tab.
   a) Select a physical volume from the **Physical Volumes** table.
   b) Select **Mirroring** and enter the names of the **Failure group 1** and **Failure group 2**. Mirroring enables you to assign physical volumes to the failure group 1 and failure group 2 that are contained by the tiers. The same data is replicated in both the failure groups. If you enable mirroring you can retrieve data when the data is lost from one failure group. To enable mirroring, you must assign the physical volumes in the table to the failure groups.

   Note: The system tier that is created in this wizard is unrestricted and it is the default tier.
8. Click **Next** or click the **Summary** tab. Verify that the shared storage pool cluster is added and complete one of the following steps:
   - Click **Back** to change the parameters.
   - Click **Finish** to add the shared storage pool cluster.

Adding tiers by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can add a tier to a Shared Storage Pool (SSP) clusters by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To add a tier to a Shared Storage Pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC).
3. Select a Shared Storage Pool cluster from the table and click **Actions** > **Add Tier**. Alternatively, you can add a tier from the cluster configuration page by clicking **Add Tier**. The **Add Tier** page opens.
4. Enter a tier name in the **Tier Name** field.
5. Enter the free threshold percentage and the overcommit threshold percentage in the **Free Threshold %** field and the **Overcommit Threshold %** field.
6. Select **Mirroring** and enter the names of the **Failure group 1** and **Failure group 2**. Mirroring enables you to assign physical volumes to the failure group 1 and failure group 2 that are contained by the tiers. The same data is replicated in both the failure groups. If you enable mirroring you can retrieve data when the data is lost from one failure group. To enable mirroring, you must assign the physical volumes in the table to the failure groups.
7. In the **Physical Volumes** table, assign **Failure group 1** and **Failure group 2** to the required physical volumes to add storage capacity.
8. Click **OK**. A tier is added to the selected Shared Storage Pool cluster.

Adding nodes by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can add a node to a Shared Storage Pool (SSP) cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To add a node to a Shared Storage Pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC).
3. Select a Shared Storage Pool cluster from the table and click **Actions** > **Add Nodes**. Alternatively, you can add a node from the cluster configuration page by clicking **Add Nodes** under the **Nodes** section. The **Add Nodes** page opens.
4. In the **Virtual I/O Server Cluster nodes** table, select the Virtual I/O Server nodes that you want to add to the Shared Storage Pool cluster.
5. Click **OK**. A cluster node is added to the selected Shared Storage Pool cluster.

Removing SSP clusters by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can remove a Shared Storage Pool (SSP) cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To remove a shared storage pool cluster that is assigned to a managed system, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC).
3. Select the cluster to be removed from the table and click **Actions** > **Remove Cluster**.

4. Click **OK** to confirm the removal of the cluster.

Changing SSP clusters by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can change a Shared Storage Pool (SSP) cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Changing the assignment of physical volumes in an SSP cluster

You can use the Hardware Management Console (HMC) to view and to change the assignment of physical volumes in a Shared Storage Pool (SSP) cluster.

Each Virtual I/O Server (VIOS) in the cluster requires at least one physical volume for the repository that is used by the Cluster Aware AIX (CAA) subsystem and one or more physical volumes for the storage pool.

When a cluster is created, you must specify one physical volume for the repository physical volume and at least one physical volume for the storage pool physical volume. The storage pool physical volumes are used to provide storage to the actual data generated by the client partitions. The repository physical volume is used to communicate with the cluster and store the cluster configuration. The maximum client storage capacity matches the total storage capacity of all storage pool physical volumes. The repository disk must have minimum 1 GB of available storage space. The physical volumes in the storage pool must have minimum 10 GB of available storage space in total.

You can use any method that is available for the storage area network (SAN) to create each physical volume with minimum 10 GB of available storage space. Map the physical volume to the partition Fibre Channel adapter for each VIOS in the cluster. The physical volumes must be mapped only to the VIOS that is connected to the shared storage pool.

After the physical volumes are allocated to a VIOS in the shared storage pool environment, the VIOS manages those physical volumes. You can change the capacity or allocation of physical volumes in a client partition.

Replacing a cluster repository disk by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can replace the assigned repository disk in a shared storage pool cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To replace the cluster repository disk in a shared storage pool cluster by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC).
3. Select a Shared Storage Pool cluster from the table and click **Actions** > **View Shared Storage Pool Cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. Under the **Repository Disk** section click **Replace Disk**. The **Replace Shared Storage Pool Repository Disk** page opens.
5. Select a cluster repository disk, from the list of those disks that are available in the table to replace the repository disk that is currently assigned to the cluster.
6. Click **OK** to apply the changes.

Removing a node by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can remove a node from a shared storage pool cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Procedure

To remove a node from a shared storage pool cluster by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC).
3. Select a Shared Storage Pool cluster from the table and click **Actions** > **View Shared Storage Pool Cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. Under the **Nodes** section click **Remove Node**.
5. Click **OK** to confirm the removal of the node.
6. Click **OK** to apply the changes.

Managing tier tasks by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can manage tier tasks in a Shared Storage Pool (SSP) cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Removing a tier

To remove a tier from a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), and the information about tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Remove Tier**.
6. Click **OK** to confirm the removal of a tier.

Removing the default tier

To remove a default tier from a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.

3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name that has *Default* as the suffix. The tier configuration page opens.
5. Click **Actions** > **Remove Default**.
6. Select another tier from the table to be the default tier.
7. Click **OK** to confirm the removal of the default tier.

Renaming a tier

To rename a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Rename Tier**. Alternatively, in the cluster configuration page, select **Actions** > **Rename Tier**. The **Rename Tier** page opens.
6. Enter a new name for the selected tier.
7. Click **OK**. The selected tier is renamed.

Setting another tier as default

To set another tier as default in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Set as Default**. Alternatively, in the cluster configuration page, select **Actions** > **Set as Default**. The **Set Default Tier** page opens.
6. Click **OK** to confirm the removal of the default tier.

Adding storage capacity

To add storage capacity to a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Add Capacity**. Alternatively, in the cluster configuration page, select **Actions** > **Add Capacity**. The **Add Capacity** page opens.
6. In the **Physical Volumes** table, assign **Failure group 1** and **Failure group 2** to the required physical volumes to add storage capacity.

   Note: Failure group 1 and Failure group 2 are displayed only if the selected tier is mirrored. If the selected tier is not mirrored, you can see *Assigned* instead of Failure group 1 and Failure group 2.
7. Click **OK**. The storage capacity is added.

Removing storage capacity

To remove storage capacity to a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Remove Capacity**. Alternatively, in the cluster configuration page, select **Actions** > **Remove Capacity**. The **Remove Capacity** page opens.
6. In the **Physical Volumes** table, unassign **Failure group 1** or **Failure group 2** from the required physical volumes to remove storage capacity.

   Note: If the selected tier is mirrored, it displays the **Failure Groups** tab. If the tier is not mirrored you can see the **Physical Volumes** tab, instead of Failure Groups tab.
7. Click **OK**. The storage capacity is removed.

Enabling mirroring

To enable mirroring in a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), and the information about tiers and nodes assigned to them.

3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Enable Mirroring**. Alternatively, in the cluster configuration page, select **Actions** > **Enable Mirroring**. The **Enable Mirroring** page opens.
6. Select the mirroring group and enter the names of **Failure group 1** or **Failure group 2** to be added. Mirroring enables you to assign physical volumes to the failure group 1 and failure group 2 that are contained by the tiers. The same data is replicated in both the failure groups. If you enable mirroring, you can retrieve data when the data is lost from one failure group. To enable mirroring, you must assign the physical volumes in the table to the failure groups.
7. Click **OK**.

Disabling mirroring

To disable mirroring in a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Disable Mirroring**. Alternatively, in the cluster configuration page, select **Actions** > **Disable Mirroring**. The **Disable Mirroring** page opens.
6. Select the mirroring group **Failure group 1** or **Failure group 2** to be removed.
7. Click **OK** to confirm the removal of the selected mirroring failure group.

Modifying thresholds

To modify threshold percentages in a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. Click **Actions** > **Modify Threshold**. Alternatively, in the cluster configuration page, select **Actions** > **Modify Threshold**. The **Modify Thresholds** page opens.
6. Enter the free threshold percentage and the overcommit threshold percentage in the **Free Threshold %** field and the **Overcommit Threshold %** field to modify the existing values.
7. Click **OK**. The threshold percentages are modified.

Restricting or unrestricting the system tier

To restrict or unrestrict the system tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers that are managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the cluster configuration page, select **Actions** > **(Un)Restrict**. The **Restrict/Unrestrict System Tier** page opens.

   Note: Restricting the system tier removes the ability to store user data on the system tier. Existing data is not affected. While, unrestricting the system tier enables user data to be stored on the system tier.
5. Click **OK** to confirm the restriction or unrestriction of the system tier.

Renaming failure groups by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can rename failure groups by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

To rename a failure group assigned to a tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), and the information about tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. In the **Failure Groups** tab, click **Rename FG**. The **Rename Failure Group** page opens.
6. Enter a failure group name in the **New Failure Group Name** field.
7. Click **OK**. The failure group is renamed.

Managing SSP physical volumes by using the All Shared Storage Pool Clusters menu

With HMC version 8.40 or later, you can manage physical volumes in a Shared Storage Pool (SSP) cluster by using the **All Shared Storage Pool Clusters** menu in the Hardware Management Console (HMC).

Replacing an SSP physical volume

To replace the existing shared storage pool (SSP) physical volume in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .

2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), and the information about tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. In the **Failure Groups** tab, click **Replace Disk**. The **Replace Shared Storage Pool Physical Volume** page opens.
6. Select the new physical volume from the table to replace the existing physical volume that is assigned to the SSP cluster. The replaced disk is free to be used for other assignments.

   Note: Ensure that at least one free physical volume, which has more size than the physical volume that is being replaced, is available.
7. Click **OK**. The physical volume is replaced.

Migrating an SSP volume to a different tier

To migrate the shared storage pool (SSP) volume to a different tier in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. In the **Shared Storage Pool Volumes** tab, click **Actions** > **Migrate to Different Tier**. The **Migrate Shared Storage Pool Volume to a different tier** page opens.
6. Select the destination tier where you want the SSP tier to be migrated. The destination tier must have enough storage space to accommodate the new tier. Depending on the size of the SSP volume, the migration might take some time to complete.

   Note: Ensure that at least one data tier or unrestricted system tier that is configured in the SSP is available, before you migrate to a different tier.
7. Click **OK**. The SSP volume is moved to a different tier.

Increasing SSP volume size

To increase the size of a shared storage pool (SSP) physical volume, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.

4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. In the **Shared Storage Pool Volumes** tab, click **Actions** > **Increase size**. The **Increase Shared Storage Pool Volume Size** page opens.
6. Enter a new storage size for the selected physical volume.
7. Click **OK**. The storage size of the selected physical volume is increased.

Removing unassigned SSP volume

To remove an unassigned shared storage pool (SSP) volume a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. In the **Shared Storage Pool Volumes** tab, click **Actions** > **Remove**. The **Remove Unassigned Shared Storage Pool Volume** page opens.
6. Click **OK** to confirm the removal of the unassigned SSP volume.

Viewing assigned partitions

To view all assigned partitions to an share storage pool (SSP) volume in a shared storage pool cluster, by using the HMC, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Shared Storage Pool Clusters**. The **All Shared Storage Pool Clusters** table is displayed. The table lists all clusters that can be accessed by the HMC (all servers managed by the HMC), with information about the tiers and nodes assigned to them.
3. Select an SSP cluster from the table and click **Actions** > **View Shared Storage Pool cluster**. Alternatively, you can click the cluster name to view the configuration details for that SSP cluster. The cluster configuration page opens.
4. In the SSP cluster table, click the tier name. The tier configuration page opens.
5. In the **Shared Storage Pool Volumes** tab, select **Show Assignment**. The partitions assigned to the SSP volumes are displayed in the table.

Managing partitions (logical partitioning)

Partitioning is the ability to make a server run as if it were two or more independent servers. When you logically partition a server, you divide the resources on the server into subsets called partitions. You can install software on a partition, and the partition runs as an independent logical server with the resources that you allocated to the partition. You can create a maximum of 1000 partitions on some servers. However, the maximum number of partitions on a server varies depending on the server configuration.

Partitions help you efficiently use system resources and increase configuration possibilities. You can use partitions to reduce the footprint of your data center by consolidating servers and maximize the use of system resources by sharing resources among multiple partitions.

You can manage the configuration of partitions and the hardware resources that are allocated to each partition by using the Manage PowerVM and Manage partition functions in the Hardware Management Console (HMC).

Note: You must activate the partition or apply the partition configuration at least once before you plan to use the Manage partition functions.

You can perform partition management functions, such as assigning processors, memory, and I/O devices to partitions, by accessing the options listed under the Properties area of the HMC graphical user interface.

You can complete most configuration updates while the partition is running.

You can run AIX, IBM i, or Linux operating systems on partitions.

Activating partitions

You can activate an IBM i, AIX, or a Linux partition by using the Hardware Management Console (HMC).

Based on which partition you want to activate, complete the steps in either the "Activating IBM i partitions" on page 60 topic or the "Activating AIX or Linux partitions" on page 61 topic. You can set the activation options to activate or network boot a partition.

Note: The partition that you choose to activate must be in the **Not Activated** state. If you select a partition that is in other states, the **Activate** option is not displayed.

Activating IBM i partitions

You can activate or network boot an IBM i partition by using the Hardware Management Console (HMC).

About this task

To activate or network boot an IBM i partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**.

   The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. To view the **Activate <IBM i partition name>** wizard, choose one of the following options:
   - In the work pane, select the partition that you want to activate and click **Actions** > **Activate**. The **Activation** wizard is displayed.
   - In the work pane, click on the partition name that you want to activate. The partition properties page opens. Click **Partition Actions** > **Operations** > **Activate**. The **Activation** wizard is displayed.
5. From the **Partition Configuration** list, select the required partition configuration profile.

   You can select only the profile that is associated to the selected partition. When you create a partition a default profile is always associated to the partition. This is indicated with the profile name being followed by **default** in parentheses.

   Note: If you choose **Current Configuration**, the **Advanced Settings** are unavailable.
6. From the **Activation Options** list, select the activation option for the partition.
   - Select **Activate** to activate the partition.

     Note: If you select **Activate**, the **Next** button is not available and you can only click **Finish** to activate and close the wizard after making all your choices in the wizard.

- Select **Network Boot** to install the operating system on the partition. Alternatively, you can select the logical partition that you want to network boot from the work pane and click **Actions** > **Netboot**. The **Network Boot** wizard is displayed. Click **Next** to configure the network settings for the logical partition.

7. Click **Advanced Settings** if you want to view and modify the following options for the selected partition:

   - **Keylock Position** establishes the power-on and power-off modes for the system. You can select the following keylock values - Do not override configuration, Manual (attended), and Normal (unattended).

     ⚠ **Attention:** The **Manual** (attended) value is not preferred value for security reasons.

   - **IPL Type** determines the copy of programs that are used by your system during initial program load (IPL).
   - **Open 5250 console** establishes a console session using the HMC 5250 emulator. This option is available only on the HMC local console and is not available on the HMC remote console.
   - **Use VSI Profile** activates the partition with Virtual Station Interface (VSI) profiles.

     Note: If the VSI attributes are not set correctly, the activation fails.

8. If you selected **Activate** from the **Activation Options** list, click **Finish** to activate the IBM i partition and close the activation wizard.
9. If you selected **Network Boot** from the **Activation Options** list, click **Next**. The **Network Settings** page opens.
10. In the **Network Settings** page, configure network adapter settings for the partition by using the following options:

    - **IPv4 or IPv6 address** to use the IPv4 or IPv6 server and client address.
    - **Boot Server IP address** to specify the IP address of the boot server that contains the network installation image for a partition. If you select **IPv4**, you must also specify the other details, such as the subnet mask and the default gateway. If you select **IPv6** you must specify the required **IPv6** settings for your system.

11. Click **Advanced Settings** to view and change the following network configuration settings for the selected partition:

    a) From the **Adapter Speed** list, select the speed of the Ethernet adapter for the target partition. By default, **Auto** is selected to enable the system to determine the required speed for the adapter. You can also select the following values - **10**, **100**, or **1000**.

    b) From the **Adapter Duplex** list, select duplex value for the Ethernet adapter. By default, **Auto** is selected to enable the system to determine the required duplex for the adapter. You can also select the **Full** or **Half** values.

    c) In the **VLAN Tag Identifier** field, specify a valid value for the virtual local area network (VLAN) tag identifier. The valid value is in the range 1 - 4094. This is an optional parameter and it is displayed only if the managed system is capable of the VLAN tagging function for the IBM i partition network boot.

12. Click **Finish** to activate with the selected network boot settings and to close the activation wizard.

Activating AIX or Linux partitions

You can activate or network boot an AIX or a Linux partition by using the Hardware Management Console (HMC).

About this task

To activate and network boot an AIX or a Linux partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**.

   The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. To view the **Activate <AIX / Linux partition name>** wizard, choose one of the following options:
   - In the work pane, select the partition that you want to activate and click **Actions** > **Activate**. The **Activation** wizard is displayed.
   - In the work pane, click the partition name that you want to activate. The partition properties page opens. Click **Partition Actions** > **Operations** > **Activate**. The **Activation** wizard is displayed.
5. From the **Partition Configuration** list, select the required partition configuration profile.

   You can select only the profile that is associated to the selected partition. When you create a partition a default profile is always associated to the partition. This is indicated with the profile name being followed by **default** in parentheses.

   Note: If you choose **Current Configuration**, the **Advanced Settings** are unavailable.
6. From the **Activation Options** list, select the activation option for the partition.
   - Select **Activate** to activate the partition.

     Note: If you select **Activate**, the **Next** button is not available. You can only click **Finish** to activate and close the wizard, after making all your choices in the current screen.
   - Select **Network Boot** to install the operating system on the partition. Alternatively, you can select the logical partition that you want to network boot from the work pane and click **Actions** > **Netboot**. The **Network Boot** wizard is displayed. Click **Next** to configure the network settings for the logical partition.
7. Click **Advanced Settings** if you want to view and modify the following options for the selected partition:
   - **Keylock Position** establishes the power-on and power-off modes that are allowed for the system. You can select the following keylock values - Do not override configuration, Manual (attended), and Normal (unattended).

     ⚠ **Attention:** The **Manual** (attended) value is not preferred value for security reasons.
   - **Boot Mode** indicates the activation type for a partition. This activation type is applicable only for AIX, Linux, or Virtual I/O Server partitions. This option is not displayed for IBM i partitions.
   - **Open vterm** opens a virtual terminal console.
   - **Use VSI Profile** activates the partition with Virtual Station Interface (VSI) profiles.

     Note: If the VSI attributes are not set correctly, the activation fails.
8. If you selected **Activate** from the **Activation Options** list, click **Finish** to activate the AIX or the Linux partition and close the activation wizard.
9. If you selected **Network Boot** from the **Activation Options** list, click **Next**. The **Network Settings** page opens.
10. In the **Network Settings** page, configure network adapter settings for the partition by using the following options:
    - **IPv4 or IPv6 address** to use the IPv4 or IPv6 server and client address.
    - **Boot Server IP address** to specify the IP address of the boot server that contains the network install image for a partition. If you select **IPv4**, you must also specify the other details, like the

subnet mask and the default gateway. If you select **IPv6** you must specify the required **IPv6** settings for your system.

11. Click **Advanced Settings** to view and change the following network configuration settings for the selected partition:

    a) From the **Adapter Speed** list, select the speed of the Ethernet adapter for the target partition. By default, **Auto** is selected to enable the system to determine the required speed for the adapter. You can also select the following values **- 10**, **100**, or **1000**.

    b) From the **Adapter Duplex** list, select duplex value for the Ethernet adapter. By default, **Auto** is selected to enable the system to determine the required duplex for the adapter. You can also select the **Full** or **Half** values.

12. Click **Finish** to activate with the selected network boot settings and to close the activation wizard.

Managing partitions

You can view and change the properties of partitions by using the Hardware Management Console (HMC).

You can view and change the following properties of a partition:

- General properties and capabilities
- Processor
- Memory
- Persistent Memory
- Physical I/O adapters

Changing partition properties and capabilities

You can view and change the partition name, view general properties of the partition, and change the virtualization capabilities by using the Hardware Management Console (HMC).

Before you begin

You can view the following general properties:

- Type, version, and IP address of the operating system.
- Machine type and serial number of the system.
- Resource configuration of a logical partition. It indicates whether all the resources that are necessary to activate the partition are available. When the **Resource Configuration** field displays **Configured**, the partition can be activated by using the current configuration. When the **Resource Configuration** field displays **Not Configured** and the partition has a last valid configuration profile, that profile is used to activate the partition. Otherwise, the partition can be activated by using a profile.

Note: When a user, who has access to the partition but does not have access to the partition profile tries to view the properties of the partition, the user is prompted to **Apply Partition Configuration**, before proceeding to view and manage the partition.

You can view or change the partition name and keylock position, add a description, and assign group tags. Additionally, if the managed system supports virtual serial number (VSN) and the managed system is not in the Enterprise Pool 2.0, you can view and manage the virtual serial number for the logical partition.

Virtualization capabilities of a partition include the following features:

Live partition mobility

Live Partition Mobility is a component of the PowerVM Enterprise Edition hardware feature that enables moving AIX, IBM i, and Linux partitions from one system to another. The mobility process transfers the system environment, including the processor state, memory, attached virtual devices, and connected users.

With the active partition mobility feature, you can move AIX, IBM i, and Linux partitions that are running, including the operating system and applications, from one system to another. The partition and the applications that run on that migrated partition do not need to be shut down.

With the inactive partition mobility feature, you can move a powered off AIX, IBM i, or Linux partition from one system to another. For more information on live partition mobility, see Partition Mobility.

You cannot migrate an IBM i partition that is configured with SR-IOV logical ports when the IBM i partition is in the **Restricted I/O Partition** mode.

You cannot migrate a logical partition that is configured with virtual Persistent Memory (PMEM) devices.

Simplified Remote Restart

When this capability is enabled, the partition state and partition configuration data is automatically stored on an HMC that manages the server. Any change to the partition configuration or profile is automatically synchronized with the data that is stored on the HMC. You can enable or disable the simplified remote restart capability only when the partition is in an inactive state.

Note: When the HMC is at Version 8.6.0, or later, and the firmware is at level FW860, or later, you can enable or disable the simplified version of the remote restart capability when the logical partition is in the `Running` state. The logical partition must not be in the `Suspended`, `Resuming`, `Migrating`, or `Remote Restarting` states.

This option is available only when the server is enabled with PowerVM Enterprise Edition and the level of firmware on your server supports the simplified remote restart capability. If a managed system is **PowerVM Partition Simplified Remote Restart Capable**, the page displays the option to manage a simplified remote restart partition only.

The partition and profile data, referred to as remote restart data, is stored on the HMC hard drive for partitions that are simplified remote restart capable. For more information on the different states of the remote restart operation, see Remote Restart State.

You cannot enable the **Simplified Remote Restart** (SRR) capability if the SR-IOV logical ports are already assigned to an IBM i partition that is in the **Restricted I/O Partition** mode.

You cannot enable the **Simplified Remote Restart** (SRR) capability if the logical partition is configured with virtual Persistent Memory (PMEM) devices.

Note: If the system key for encrypting the platform keystore data is not a user-defined key, then you cannot create a logical partition with both the platform keystore and simplified remote restart (SRR) capabilities enabled.

About this task

To view and change the properties and capabilities of the partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** area, click **Properties** > **General Properties** to view and change the properties of the selected partition.
5. Enter a name in the **Partition Name** field to change the name of the partition.
6. Select the **Key Lock Position** to be either **Manual** or **Normal**.
7. Enter an optional description in the **Description** field to further identify the logical partition.

8. In the **Group Tags** field, select from the list of available tag assignments for the groups to which the partition belongs. If the partition does not belong to any group, the group tags list is empty.
9. If the managed system supports virtual serial number (VSN) and the managed system is not in the Enterprise Pool 2.0, you can view and manage the virtual serial number for the logical partition in the **Virtual Serial Number** field as described in the following sections:

   a) When the logical partition is in **Running** state, you can only view the virtual serial number. If the virtual serial number is assigned to the logical partition, the **Virtual Serial Number** field displays the virtual serial number. **No VSN** is displayed in the **Virtual Serial Number** field if the virtual serial number is not assigned to the logical partition.

   b) When the logical partition is in **Not activated** state, you can view and modify the virtual serial number.

      - If the virtual serial number is assigned to the logical partition, the **Virtual Serial Number** field displays the virtual serial number. Click and enable the **No VSN** check box if you do not want to assign the virtual serial number to the logical partition.
      - If the virtual serial number is not assigned to the logical partition, the **Virtual Serial Number** field displays the following options:
        - **No VSN**: Select **No VSN** option if you do not want to assign the virtual serial number to the logical partition.
        - **Auto-assign**: Select **Auto-assign** option if you want the system to automatically assign a virtual serial number to the logical partition.
        - **Select from pool**: Choose the **Select from pool** option if you want to manually assign a virtual serial number to the logical partition. Click **Select VSN** to open the **Virtual Serial Numbers** window. The window lists the virtual serial number groups and the available virtual serial number that can be assigned to the logical partition. Select a virtual serial number from the list. Click **OK** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the window.

   Note: When the Power® firmware is at level FW950 and if the managed system already have logical partitions to which the virtual serial number is assigned, the logical partition cannot be added to the Enterprise Pool 2.0. Alternatively, if the managed system is already in the Enterprise Pool 2.0, then the managed system cannot assign a virtual serial number to the logical partition.
10. Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the page.

Disabling Live Partition Mobility

You can disable the Live Partition Mobility feature of a logical partition by using the Hardware Management Console (HMC).

About this task

The HMC provides the **Disable Migration** option to disable the Live Partition Mobility feature at a logical partition level. This option can be used by customers to address application licensing requirements of Independent Software Vendors (ISV). To disable the Live Partition Mobility feature for a logical partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to disable the Live Partition Mobility feature and click **Actions** > **View Partition Properties**.
4. In the **Properties** area, click **Properties** > **General Properties**.

5. Click the **Advanced** tab and in the **Advanced settings** area, select the **Disable Migration** check box.
6. Click **Save**.

   Some Independent Software Vendors might require you to purchase a license for all systems where their application can be migrated. Rather than disabling the Live Partition Mobility feature at a system level, IBM provides this logical partition level mechanism that can be audited, to disable migration to address ISV licensing requirements while preserving the ability to leverage migration for applications running on other logical partitions on the system.

   Note: IBM software does not stipulate such licensing requirements.

   The **Disable Migration** option is supported on all firmware versions, and when the system is managed by an HMC that is at version 8.4.0, or later. Also, you can run the **chsyscfg** command with a value 1 for the *migration_disabled* attribute, from the HMC command line. To disable the Live Partition Mobility feature of a logical partition during partition creation, run the **mksyscfg** command with a value 1 for the *migration_disabled* attribute, from the HMC command line. The **Disable Migration** option is also supported by the Representational State Transfer (REST) application programming interfaces (APIs).

   Note: PowerVM NovaLink supports the **Disable Migration** option when the system is co-managed by an HMC. However PowerVM NovaLink does not provide an option to disable Live Partition Mobility feature.

Viewing system event logs for the Live Partition Mobility disable operation

Any changes that are made to the **Disable Migration** option provided by the Hardware Management Console (HMC) is logged as a system event, and can be checked for auditing purposes. A system event is also logged when the remote restart or simplified remote restart capability is set. The system event logs are read only and cannot be modified.

A system event is logged when the following actions occur:

- The remote restart, simplified remote restart, or Live Partition Mobility attributes are set during creation of a logical partition.
- The remote restart, simplified remote restart, or Live Partition Mobility attributes are changed.
- When you restore profile data. For more information about restoring profile data, see Restoring profile data.

You can view the system events by running the **lssvcevents** command from the HMC command-line interface. You can also view the system events by using the graphical user interface (GUI). For more information about using the GUI, see Console Events Logs. By running the **chhmc** command from the HMC command-line interface, these system events can also be sent to a remote server that is on the same network as the HMC.

The following system events can be logged:

Table 2. Event ID and the corresponding message string

Event ID	Event message String
2420	User name {0}: Disabled partition migration for partition {1} with ID {2} on managed system {3} with MTMS{4}
2421	User name {0}: Enabled partition migration for partition {1} with ID {2} on managed system {3} with MTMS{4}
2422	User name {0}: Disabled Simplified Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}
2423	User name {0}: Enabled Simplified Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}
2424	User name {0}: Disabled Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}

Table 2. Event ID and the corresponding message string (continued)

Event ID	Event message String
2425	User name {0}: Enabled Remote Restart for partition {1} with ID {2} on managed system {3} with MTMS{4}

The following are examples of System Events:

- Command to check which logical partitions managed by an HMC have the Live Partition Mobility feature disabled or enabled:

  ```
  lssvcevents -t console | grep vclient
  ```

  The following examples show a sample output of the **lssvcevents** command:

  ```
  time=10/30/2015 10:11:32,text=HSCE2521 UserName hscroot: Enabled partition migration for
  partition
  vclient10 with Id 10 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

  ```
  time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for
  partition
  vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

- Command to check which logical partitions managed by an HMC have the Live Partition Mobility feature disabled:

  ```
  lssvcevents -t console | grep HSCE2520
  ```

  The following example shows a sample output of the **lssvcevents** command:

  ```
  time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for
  partition
  vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

- Command to check which logical partitions managed by an HMC have the Live Partition Mobility feature disabled or enabled for a particular system (1234567):

  ```
  lssvcevents -t console | grep  "partition migration for partition" | grep 1234567
  ```

  The following examples show a sample output of the **lssvcevents** command:

  ```
  time=10/30/2015 10:11:32,text=HSCE2521 UserName hscroot:  Enabled partition migration for
  partition
  vclient10 with Id 10 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

  ```
  time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for
  partition
  vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567.
  ```

- Command to check whether a specific logical partition (vclient9) in a specific system (1234567) managed by an HMC has the Live Partition Mobility feature disabled or enabled:

  ```
  lssvcevents -t console | grep  "partition migration for partition vclient9" | grep 1234567
  ```

  The following example shows a sample output of the **lssvcevents** command:

  ```
  time=10/30/2015 10:01:35,text=HSCE2520 UserName hscroot: Disabled partition migration for
  partition
  vclient9 with Id 9 on Managed system ct05 with MTMS 8205-E6D*1234567
  ```

Changing advanced partition settings

You can view and change advanced settings of a partition by using the Hardware Management Console (HMC).

Before you begin

The advanced settings for a partition include the following options:

Enable connection monitoring
Monitors the connection between the partition and the HMC.

Enable redundant error path information
If you enable the redundant error-path reporting, the partition reports common server hardware errors and partition hardware errors to the HMC. If you disable redundant error-path reporting, the partition reports only partition hardware errors to the HMC. If you want to move a partition, disable the redundant error-path reporting.

Enable time reference
Synchronize the PowerVM Hypervisor and Service Processors time-of-day based on the time-of-day setting of the concerned partition and other Time Reference partitions.

Disable migration
You can disable the Live Partition Mobility feature for an AIX, Linux, or IBM i partition.

Service partition

Indicates whether the partition is the service partition for the managed system. The service partition is the IBM i logical partition on an IBM System i® managed system that you can configure to apply server firmware updates to the service processor or to the Hypervisor and to communicate server common hardware errors to IBM. These abilities are useful if the HMC is undergoing maintenance or is otherwise unable to perform these functions. You must change the service partition on the managed system through the managed system properties.

Enable virtualized trusted platform module (VTPM)
With the HMC Version 7 Release 7.4.0, or later, and IBM POWER7® processor-based servers with firmware at level 7.4, or later, you can enable the virtual trusted platform module (VTPM) on an AIX or Linux partition. A partition that is VTPM enabled supports the Trusted Boot capability. Trusted Boot is a capability that is supported on the Power Security and Compliance (PowerSC) Standard Edition. Up to 60 partitions per server can be configured to have their own unique VTPM by using the HMC. The VTPM is used to record the system boot and, in association with the AIX Trusted Execution technology, provides security and assurance of the boot image on disk, on the entire operating system, and in the application layers.

Tagged I/O Settings

You can view, configure and specify the exact I/O devices that you want to use for a logical partition to perform specific functions.

Note: You cannot activate a logical partition if the SR-IOV logical ports are configured as a primary load source or an alternate load source in the **Tagged I/O Settings**.

Enable performance information collection
Enable the operating system on a partition to collect performance information.

Restricted I/O Mode SR-IOV Logical Port Assignment Capable
When the HMC is at version V9.1.940, or later, when the firmware is at level FW940, or later, and when the partition is running the latest version of IBM i, the HMC supports assignment of SR-IOV logical ports to IBM i partitions that are in the **Restricted I/O Partition** mode. In the **Advanced** settings area, the **Restricted I/O Mode SR-IOV Logical Port Assignment Capable** field is displayed.

The **Restricted I/O Mode SR-IOV Logical Port Assignment Capable** field displays the following values:

- **Supported**: The SR-IOV logical port can be assigned to the IBM i partitions that are in the **Restricted I/O Partition** mode.

- **Unsupported**: The SR-IOV logical port cannot be assigned to the IBM i partitions that are in the **Restricted I/O Partition** mode.
- **Unavailable**: No information is available because the IBM i partition has never been activated.

Note: The **Restricted I/O Mode SR-IOV Logical Port Assignment Capable** field is displayed only for the IBM i partitions. The IBM i partition must be activated at least once to know the exact value in the **Restricted I/O Mode SR-IOV Logical Port Assignment Capable** field.

Restricted I/O partition

Determines whether an IBM i partition can be migrated by using the Live Partition Mobility (LPM) feature. You can migrate the IBM i partition only if you select the **Restricted I/O Partition** option. On servers that do not support the IBM i partitions with native I/O capability, you must always enable this option. On servers running firmware level FW860 or later, the IBM i native I/O capability of the server is available on the **Licensed Capabilities** page. This option can be enabled only when the partition is stopped.

Note: The Restricted I/O partition setting is applicable only to IBM i partitions.

OptiConnect

A feature of the IBM i operating system that allows a user to connect multiple System i systems by using SPD bus, high-speed link (HSL) loop, or virtual interpartition technologies. This option can be enabled only when the partition is stopped.

Enable Electronically report errors that cause partition termination or require attention

Select this option to set the HMC to send an electronic report to service and support whenever this IBM i logical partition terminates abnormally or encounters an error that requires service. (The HMC does not report errors that requires user intervention.) Use this feature to enable automatic service calls for IBM i logical partitions that run mission-critical applications. This field displays only for IBM i logical partitions.

Platform KeyStore

When you are using HMC version 9.2.950, or later, and when the firmware is at level FW950, you can view and manage the platform keystore size that is used to save the partition key data. When the managed system is in **stand-by** or **operating** state, you can view and specify the platform keystore size. You cannot decrease the platform keystore size for a logical partition. If the platform keystore does not contain any data, you can specify the value **0** in the **KeyStore Size** field to disable the platform keystore feature. You can enable or increase the platform keystore size only when the logical partition is in the **Not activated** state.

Note: When the logical partition is enabled with the platform keystore and simplified remote restart (SRR) features, and if you want to change the system key from a user-defined key to a default system key, then you must manually disable either the platform keystore or the SRR feature on the logical partition.

Note: If the system key for encrypting the platform keystore data is not a user-defined key, then you cannot create a logical partition with both the platform keystore and simplified remote restart (SRR) capabilities enabled.

Supported hardware accelerator types

The **Supported Hardware Accelerator Types** table is displayed only if the managed system has user-mode hardware accelerator for PowerVM®. You can specify the GZIP Quality of Service (QoS) credits for a logical partition. The HMC sets the QoS, if sufficient credits are available for the managed system. The logical partition uses these credits to access the shared hardware accelerators. The following are the prerequisites for a logical partition to receive QoS credits.

- The managed system supports hardware accelerator enablement.
- The managed system supports the GZIP hardware accelerator type.
- QoS enablement must be supported at the operating system level.
- The managed system has sufficient hardware accelerator QoS credits to assign to the partition.
- The logical partition must be either an AIX, Linux, or a VIOS partition.

About this task

To view and change the advanced settings of the partition, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** area, click **Properties** > **General Properties** to view and change the properties of the selected partition.
5. Click the **Advanced** tab. The **Advanced Settings** options are displayed.
6. To enable the advanced settings in the selected partition, select the following options:
   a) **Enable Connection Monitoring** to monitor the connection.
   b) **Enable Redundant Error Path Reporting** to report common server hardware errors and partition hardware errors.
   c) **Enable Time Reference**
   d) **Service Partition**
   e) **Disable Migration** to disable the Live Partition Mobility feature for an AIX, Linux, or IBM i partition.
   f) **Restricted I/O Partition**
   g) Specify a value in the **Maximum Virtual Adapters** field.
   h) **Enable Virtualized Trusted Platform Module (VTPM)** to record the system boot and to provide security and assurance of the boot image on disk, on the entire operating system, and in the application layers.
   i) **Tagged I/O Settings** - In the **Advance** settings area, click the **Tagged I/O Settings** tab if you want to view or configure the I/O devices that you want to use for a logical partition. The **Tagged I/O Device Details** window is displayed.
      - Select the load source that must be used by the system to start the logical partition from the **Load Source** list.
      - Select an alternate device from the **Alternative Restart Device** list.
      - Select the console from the **Console** list.
      - Select an alternate console from the **Alternate Console** list.
      - Select the operations console from the **Operations Console** list.
      - Click **OK** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the window.
   j) **Enable Performance Information Collection**
   k) **Enable electronically report errors that cause partition termination or require attention**
7. From the **Save configuration changes** list, select one of the following options:
   a) **Enabled** to apply and save the settings that you made on the partition.
   b) **Disabled** to cancel the settings that you made on the partition.
   c) **Disabled until next activate or apply** to temporarily disable the settings that you made and to apply the settings later when you activate the partition.
8. In the **Supported Hardware Accelerator Types** table, specify a value in the **QoS** field.
9. From the **Secure Boot** list, select an option. The options are **Disabled**, **Enabled and Log only**, or **Enabled and Enforced**.

10. Specify a value in the **KeyStore Size** field.
11. Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the page.

Validating the configuration of a logical partition before the migration operation

You can use the validation function on the Hardware Management Console (HMC) to validate the configuration of a logical partition before starting the migration operation. If you are using HMC V9.1.940, or later, and when the firmware is at level FW940, or later, you can validate a logical partition that is configured with Single Root I/O Virtualization (SR-IOV) logical ports.

About this task

To validate the configuration of a logical partition before starting the migration operation, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition.
3. Click **View System Partitions**. The **All Partitions** page is displayed.
4. In the work pane, select the logical partition and click **Actions** > **Mobility** > **Validate**. The **Partition Migration Validation** window is displayed.
5. Follow the steps in the **Partition Migration Validation** window.
6. You can use the **Only migrate SR-IOV logical ports when possible** option to validate or migrate a logical partition that is configured with SR-IOV logical ports.

   - When the **Only migrate SR-IOV logical ports when possible** check box is enabled, the HMC performs the following task:

     – The HMC fails the validation or migration operation if the HMC cannot find an SR-IOV physical port to re-create any of the source SR-IOV logical ports that are migratable before the client logical partition is moved to the destination system.

     – If the client logical partition is already moved to the destination system and if the re-creation of the migratable SR-IOV logical ports on the destination system fails, the HMC does not move the client logical partition back to the source system. You can dynamically add the migratable SR-IOV logical ports on the client logical partition after the migration operation completes.

   - When the **Only migrate SR-IOV logical ports when possible** check box is disabled, the HMC performs the following task:

     – The HMC does not fail the validation or migration operation even if the HMC cannot find an SR-IOV physical port to re-create any of the migratable SR-IOV logical ports on the source system. Also, the HMC does not attempt to re-create those SR-IOV logical ports. You can dynamically add the migratable SR-IOV logical ports on the client logical partition after the migration operation completes.

   Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.
7. Complete the steps in the **Partition Migration Validation** window, and click **Validate**.

8. If the validation operation fails, the **Validation Errors/Warnings** window is displayed that helps you to troubleshoot the configuration problems.
   - Click **All Messages** to view the warnings and error messages along with the verbose log information about the validation operation. The log information also includes data about successful and failed steps.
   - Click **Step Results** to view the log information about the validation operation. The log information also includes data about successful and failed steps.

Migrating a logical partition

You can use the migration function on the Hardware Management Console (HMC) to migrate a logical partition. If you are using HMC V9.1.930, or earlier, you cannot migrate a logical partition that is configured with Single Root I/O Virtualization (SR-IOV) logical ports. If you are using HMC V9.1.940, or later, and when the firmware is at level FW940, or later, you can migrate a logical partition that is configured with SR-IOV logical ports.

About this task

To migrate a logical partition from one server to another server by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition.
3. Click **View System Partitions**. The **All Partitions** page is displayed.
4. In the work pane, select the logical partition and click **Actions** > **Mobility** > **Migrate**. The **Partition Migration** window is displayed.
5. Follow the steps in the **Partition Migration** window.
6. In the **Migration Information** section, you can use the **Only migrate SR-IOV logical ports when possible** option to migrate a logical partition that is configured with SR-IOV logical ports.
   - When the **Only migrate SR-IOV logical ports when possible** check box is enabled, the HMC performs the following task:
     - The HMC fails the migration operation if the HMC cannot find an SR-IOV physical port to re-create any of the source SR-IOV logical ports that are migratable before the client logical partition is moved to the destination system.
     - If the client logical partition is already moved to the destination system and if the re-creation of the migratable SR-IOV logical ports on the destination system fails, the HMC does not move the client logical partition back to the source system. You can dynamically add the migratable SR-IOV logical ports on the client logical partition after the migration operation completes.
   - When the **Only migrate SR-IOV logical ports when possible** check box is disabled, the HMC performs the following task:
     - The HMC does not fail the migration operation even if the HMC cannot find an SR-IOV physical port to re-create any of the migratable SR-IOV logical ports on the source system. Also, the HMC does not attempt to re-create those SR-IOV logical ports. You can dynamically add the migratable SR-IOV logical ports on the client logical partition after the migration operation completes.

   Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

7. Complete the steps in the **Partition Migration** window, and click **Finish**.
8. In the **Validation Errors/Warnings** section, you can view the error message information that helps you to troubleshoot the configuration problems.

    - Click **All Messages** to view the warnings and error messages along with the verbose log information about the migration operation. The log information also includes data about successful and failed steps.
    - Click **Step Results** to view the log information about the migration operation. The log information also includes data about successful and failed steps.

Changing processor settings

You can view and change the settings of the shared and dedicated processors that are assigned to a partition by using the Hardware Management Console (HMC).

Before you begin

You can change the number of virtual processors and processing units that are assigned to the partition. The views and controls that are displayed depend on whether the processor is dedicated or shared, or running or stopped.

You can set a partition to use either processors that are dedicated to the partition or processors that are shared with other partitions. If a partition uses dedicated processors, you must allocate processors (in increments of whole numbers) to the partition. A partition that uses dedicated processors cannot use any processing capacity beyond the processors that are assigned to the partition.

By default, all physical processors that are not dedicated to specific partitions are grouped in a shared processor pool. You can allocate a specific amount of the processing capacity in this shared processor pool to each partition that uses shared processors. With some models, you can use the HMC to configure multiple shared processor pools. These models have a default shared processor pool that contains all the processor resources that do not belong to partitions that use dedicated processors or partitions that use other shared processor pools. The other shared processor pools on these models can be configured with a maximum processing unit value and a reserved processing unit value. The maximum processing unit value limits the total number of processors that can be used by the partitions in the shared processor pool. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped partitions within the shared processor pool.

You can set a partition that uses shared processors to use a minimum of 0.10 processing units, which are approximately a 10th of the processing capacity of a single processor. When the firmware is at level 7.6, or later, you can set a partition that uses shared processors to use a minimum of 0.05 processing units, which are approximately a 20th of the processing capacity of a single processor. You can specify the number of processing units to be used by a shared processor partition to the 100th of a processing unit. In addition, you can set a shared processor partition such that, if the partition requires more processing capacity than its assigned number of processing units, the partition can use processor resources that are not assigned to any partition or processor resources that are assigned to another partition but that are not being used by the other partition. Some server models might require you to enter an activation code before you can create partitions that use shared processors.

If the operating system and server model supports, you can allocate up to the entire processing capacity on the managed system to a single partition. You can configure your managed system such that it does not comply with the software license agreement for your managed system. However, if you operate the managed system in such a configuration, you might receive out-of-compliance messages.

Shared processors are physical processors that share processing capacity among multiple partitions. The ability to divide physical processors and share them among multiple partitions is known as the Micro-Partitioning® technology.

Partitions that use shared processors can have a sharing mode of capped or uncapped. An uncapped partition is a partition that can use more processor power than its assigned processing capacity. The amount of processing capacity that an uncapped partition can use is limited only by the number of virtual processors that are assigned to the partition or the maximum processing units that are allowed by the shared processor pool that the partition uses. In contrast, a capped partition is a partition that cannot use more processor power than its assigned processing units.

Dedicated processors are whole processors that are assigned to a single partition. If you choose to allocate dedicated processors to a partition, you must allocate at least one processor to that partition. Likewise, if you choose to remove processor resources from a dedicated partition, you must remove at least one processor from the partition. On systems that are managed by an HMC, dedicated processors are assigned to partitions that use partition profiles.

A virtual processor is a representation of a physical processor core to the operating system of a partition that uses shared processors.

About this task

To view and change the settings of the processor, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** area, click **General Properties** to view and change the properties of the selected partition.
5. In the **Properties** area, click **Processors** to view the shared and dedicated processors.
6. Select a processor mode that is assigned to the selected partition:
   - When the partition is in the running state, and the processor is set to the **Dedicated** mode, complete the following steps:
     a. You can enter a value or adjust the **Processors** tab for the number of processors that is assigned to the partition.
     b. Click **Advanced** to change the advanced processor settings for the partition.
   - When the partition is in the not activated state, and the processor is set to the **Dedicated** mode, complete the following steps:
     a. From the **Processor Mode** list, change the mode of the processor to shared or dedicated.
     b. Enter values or adjust the **Processors** tab for the maximum, assigned, and minimum number of dedicated processors for the partition.
     c. From the **Processor Compatibility Mode** list, select the compatibility mode of the processor.
     d. Select the **Idle Processing Sharing** check box to enable and use the idle processors that belong to the powered-off shared partition.
   - When the partition is in the running state, and the processor is set to the **Shared** mode, complete the following steps:
     a. Enter a value, or adjust the **Virtual Processors** bar and the **Processing Units** bar for the assigned number of virtual processors and processing units for the partitions in the shared processor pool.
     b. Adjust the capped and uncapped setting for the partition in the shared processor pool.

- When the partition is in the not activated state, and the processor is set to the **Shared** mode, complete the following steps:

   a. From the **Processor Mode** list, select an option to change the mode of the processor to shared or dedicated.
   b. From the **Shared Processor Pool** list, select an available pool to change the shared processor pool.
   c. Adjust the capped and uncapped tab setting for the partition in the shared processor pool.
   d. Enter values or adjust the **Virtual Processors** tab for the maximum, assigned, and minimum number of shared processors for the partition.
   e. From the **Processor Compatibility Mode** list, select the compatibility mode of the processor.

7. Optional: Enter a timeout value in the **Timeout** field.
8. Optional: Click the **Force** option to force a particular operation to be performed on the logical partition immediately.
9. Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

Changing memory settings

You can view and change the settings of the shared and dedicated memory that is assigned to a partition by using the Hardware Management Console (HMC).

Before you begin

You can change the memory that is allocated to the partition. The views and controls that are presented depend on whether the memory is dedicated or shared, an whether the partition is or running or stopped.

Processors use memory to temporarily hold information. Memory requirements for partitions depend on the partition configuration, I/O resources that are assigned, and applications used.

Memory can be assigned in increments of 16 MB, 32 MB, 64 MB, 128 MB, and 256 MB. The default memory block size varies according to the amount of configurable memory in the system. On systems that are managed by an HMC, memory is assigned to partitions using partition profiles.

Dedicated memory is physical system memory that you allocate to a partition that uses dedicated memory and is reserved for use by the dedicated memory partition until you remove the memory from the dedicated memory partition or delete the dedicated memory partition.

Depending on the overall memory in your system and the maximum memory values you choose for each partition, the server firmware must have enough memory to complete partition tasks. The following factors influence server firmware memory requirements:

- Number of dedicated memory partitions
- Partition environments of the dedicated memory partitions
- Number of physical and virtual I/O devices that are used by the dedicated memory partitions
- Maximum memory values allocated to the dedicated memory partitions

Note: Firmware level updates can also change the server firmware memory requirements. Larger memory block sizes can exaggerate the memory requirement change.

When you select the maximum memory values for each dedicated memory partition, consider the following points:

- Maximum values affect the hardware page table (HPT) size for each dedicated memory partition
- The logical memory map size for each dedicated memory partition

If the server firmware detects that a memory module failed or is about to fail, the server firmware creates a serviceable event. The server firmware can also unconfigure the failing memory module automatically, depending on the type of failure and the deconfiguration policies that you set up by using the Advanced System Management Interface (ASMI). You can also unconfigure a failing memory module manually by using the ASMI. If a memory module failure causes the entire managed system to shut down, the managed system restarts automatically if the managed system is in normal initial program load (IPL) mode. When the managed system restarts itself, or when you restart the managed system manually, the managed system attempts to start the dedicated memory partitions that were running at the time of the memory module failure with their minimum memory values. If the managed system does not have enough memory to start all of the dedicated memory partitions with their minimum memory values, the managed system starts as many dedicated memory partitions as it can with their minimum memory values. After the managed system starts the maximum possible number of dedicated memory partitions, the managed system distributes the leftover memory resources among the running dedicated memory partitions, in proportion to the required memory values of the dedicated memory partitions.

Using *Huge pages* can improve performance in specific environments that require a high degree of parallelism, such as in the DB2® database. You can specify huge-page memory that can be used for the shared-memory buffer pools in the DB2 database. For logically partitioned systems, you can specify the minimum, wanted, and maximum number of huge pages to allocate to a partition when you create the partition or partition profile.

On managed systems that support huge-page memory, you can use the HMC to set the value for the huge-page memory pool. You can also specify values for the number of huge pages to allocate to partitions.

To use huge-page memory, you must ensure that your system has adequate memory resources to dedicate to the huge-page memory pool. The huge-page memory pool is a region of the system memory that is mapped as 16 GB page segments and is managed separately from the base memory of the system.

About this task

To view and change the settings of the memory, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** area, click the **Memory** tab to view the properties of the running logical partition that is using the dedicated or the shared memory.
5. Select a memory mode that is assigned to the selected partition:
   - When the partition is in the running state, and the memory is set to the **Dedicated** mode, complete the following steps:
     a. You can enter a value or adjust the **Memory Allocation** tab for the assigned memory that is assigned to the partition.
     b. Click **Advanced** to view the advanced memory settings for the partition.
   - When the partition is in the not activated state, and the memory is set to the **Dedicated** mode, complete the following steps:
     a. You can enter a value or adjust the values for **Memory Allocation** tab for the maximum, assigned, and minimum memory that is assigned to the partition.
     b. Click **Advanced** to change the advanced memory settings for the partition.
     c. Select the **Enable Active Memory Expansion** check box to enable the active memory expansion feature for the partition.
     d. Enter a value for the **Active Memory Expansion** field. The value must be in the range 1.0 - 10.0.
     e. Select the **Huge Page Memory** check box to enable the huge-page memory feature for the partition.

   f. Enter values for the **Minimum**, **Assigned**, and **Maximum** fields.
   g. Select the **BSR Array** check box to allocate barrier-synchronization register (BSR) arrays to the partition.
   h. Enter values for the **Total**, **Assigned**, and **Available** fields.
   i. From the **Memory Mode** list select shared to set the mode to shared. You can change the memory mode to shared only when there is a share memory pool is available. Also, you can change the memory mode to shared only when the processor is also set to the shared mode.

   Note: BSR is not supported on POWER8 processor-based servers.

- When the partition is in the running state, and the memory is set to the **Shared** mode, complete the following steps:
   a. You can enter a value or adjust the **Memory Allocation** tab for the assigned memory that is assigned to the partition.
   b. Click **Avanced** to change the advanced memory settings for the partition.
   c. From the **Assigned I/O Entitled Memory** option, select **Auto** or **Manual**.
   d. Enter values for the **Assigned I/O Entitled Memory** and **Memory Weight** fields.
- When the partition is in the not activated state, and the memory is set to the **Shared** mode, complete the following steps:
   a. Change the mode of the memory to shared or dedicated.
   b. Enter a value, or adjust the **Memory Allocation** tab for the maximum, assigned, and minimum dedicated memory that is assigned to the partition.
   c. Click **Avanced** to change the advanced memory settings for the partition.
   d. From the **Assigned I/O Entitled Memory** option, select **Auto** or **Manual**. When you select the manual option, you must also enter values for the **Assigned I/O Entitled Memory** and **Memory Weight** fields.
   e. From the **Memory Mode** list select dedicated to set the mode to dedicated.

6. Optional: Enter a timeout value in the **Timeout** field.
7. Optional: Click the **Force** option to force a particular operation to be performed on the logical partition immediately.
8. Click **Save** to apply the changes. Alternatively, click **Cancel** to reject the changes and close the page.

Managing persistent memory volume

You can view and manage the persistent memory volume that is assigned to a logical partition by using the Hardware Management Console (HMC).

Before you begin

When you are using HMC V9.1.940, or later, and when the firmware is at level FW940, or later, you can view and manage the persistent memory volume that is assigned to a logical partition by using the Hardware Management Console (HMC). You can manage the persistent memory volume only when the logical partition is in the **Not activated** state. You cannot add virtual Persistent Memory (PMEM) volumes to a logical partition in which the **Simplified Remote Restart** (SRR) capability is enabled.

About this task

To view and manage the persistent memory volume, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .

2. Click **All Partitions**. The **All Partitions** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties, and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** area, click the **Persistent Memory** tab to view the properties of the logical partition that is using persistent memory volume.
5. Click **Add** to create the persistent memory volume. The **Add Volume** window is displayed.
   a) Specify a name for the persistent memory volume in the **Volume Name** field.
   b) Specify a value for the size of the persistent memory volume in the **Volume Size** field.
   c) Select the **Affinity** check box if you want the operating system to get information about the amount of memory that is allocated across multiple DIMMs.
   d) Click **OK** to create the persistent memory volume. Alternatively, click **Cancel** to reject the changes and to close the window.
6. In the **Persistent Memory** page, select an existing persistent memory volume. Click **Edit** if you want to change the name of the persistent memory volume.
7. In the **Persistent Memory** page, select an existing persistent memory volume. Click **Remove** if you want to delete the persistent memory volume.

Managing physical I/O adapters

You can view and change physical I/O adapter assignment for a partition by using the Hardware Management Console (HMC).

You can dynamically add, remove, and move physical I/O devices and slots, to and from running partitions, by using an HMC. You can share infrequently used I/O devices, such as optical disk drives among multiple partitions.

You can specify that I/O devices or slots are required for a partition. If you specify that an I/O device or slot is not required, the I/O device or slot can be shared with other partitions, or the I/O device or slot is optional. If you specify that an I/O device or slot is required (or dedicated) and if the I/O device or slot is unavailable, or is in use by another partition, you cannot activate the partition.

Note: If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. All configuration changes are lost when the partition profile is next activated. If you want to save your new partition configuration, either change the partition profile or save the partition configuration to a new partition profile.

You can specify a timeout value in the **Timeout** field. You can also click the **Force** option to force a particular operation to be performed on the logical partition immediately.

Adding a physical I/O adapter to a partition

You can dynamically add a physical I/O slot, the adapter, and devices that are connected to the slot to an active partition by using the Hardware Management Console (HMC). You can add I/O capabilities to an active partition without shutting down the partition.

Before you begin

Consider the following conditions when you add physical I/O slots to a Linux partition:

- A Linux distribution that supports dynamic partitioning is installed on the Linux partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER® systems website.

You cannot add physical I/O devices and slots to partitions that use shared memory. You can allocate only virtual adapters to partitions that use shared memory.

About this task

To dynamically add a physical I/O adapter to an active partition by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** area, click **Physical I/O Adapters**. The table lists all the adapters that are available for the partition.
5. Click **Add Adapter**. The **Add Physical I/O Adapter(s)** page opens.
6. Select an I/O adapter from the **Add Physical I/O Adapter(s)** list that you want to add to the partition. You can view the adapters that are available in the other drawers of the server by clicking the **View** list. You can also narrow your search for available adapters, by using the filter to list adapters based on physical location code.
7. Click **Save** after you select the I/O adapter. Alternatively, click **Cancel** to reject the changes and close the page.

Removing a physical I/O adapter from a partition

You can dynamically remove a physical I/O slot, the adapter, and devices that are connected to that slot by using the Hardware Management Console (HMC). You can reassign the physical I/O adapter to other partitions.

Before you begin

Ensure that the devices that are attached to the managed system through the physical I/O slot that you want to remove are not running, by using the operating system commands.

⚠ **Attention:** The dynamic removal of a physical I/O slot that controls disk drives can cause unpredictable results, such as partition failure or loss of data.

Consider the following conditions when you remove a physical I/O slot from a Linux partition:

- A Linux distribution that supports dynamic partitioning is installed on the Linux partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux partition. To download the DynamicRM tool package, see the Service and productivity tools for Linux on POWER systems website.

About this task

To dynamically remove a physical I/O adapter from an active partition by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.

4. In the **Properties** area, click **Physical I/O Adapters**.
5. From the table that lists the assigned physical adapters, right-click the physical adapter that you want to remove and select **Remove Adapter**.
6. Click **Save** after you select the I/O adapter. Alternatively, click **Cancel** to reject the changes and close the page.

Results

The selected physical I/O adapter is removed from partition.

Managing partition profiles for logical partitions

You can manage the partition profiles for your logical partitions using the Hardware Management Console (HMC). You can change the resource specifications stored in your partition profiles as your needs change.

Creating a partition profile

You can create a new partition profile using the Hardware Management Console (HMC).

About this task

To create a new partition profile using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. In the **Manage Profiles** wizard, click **Actions** > **New**. The **Create Lpar** wizard is displayed.
5. In the **Partition Profile** page, specify a name for the logical partition profile in the **Profile name** field.
   a) Select the **Use all the resources in the system** check box if you want the logical partition to have all the resources that are available in the system.
   b) Click **Next**. Alternatively, click **Cancel** to reject the changes and to close the window.
6. In the **Processors** page, select **Shared** to assign partial processor units from the shared processor pool or select **Dedicated** to assign the entire processor that can only be used by the logical partition, and click **Next**.
7. In the **Processing Settings** page, specify values for the processor units in the **Minimum processing units**, **Desired processing units**, and **Maximum processing units** fields. Select a shared processor pool from the **Shared processor pool** list.
   a) In the **Virtual Processors** section, specify values for the virtual processor in the **Minimum virtual processors**, **Desired virtual processors**, and **Maximum virtual processors** fields.
   b) If you want to specify uncapped weight value for the virtual processors, select the **Uncapped Weight** check box and specify a value in the **Uncapped Weight** field, and click **Next**.
8. In the **Memory Settings** page, specify values for the logical partition memory in the **Minimum Memory**, **Desired Memory**, and **Maximum Memory** fields, and click **Next**.
9. In the **I/O** page, select the adapters that you want to be included in the partition profile from the adapter list and click **Add as required**, or **Add as desired**. Click **Remove** to remove the selected adapter from the partition profile, and click **Next**.
10. In the **Virtual Adapters** page, select the virtual adapter from the available adapter list and click **Actions** > **Properties** to view the properties of the virtual adapter.
    a) In the **Maximum virtual adapters** field, specify the number of maximum virtual adapters for the logical partition.

   b) To create a virtual adapter, click **Actions** > **Create Virtual Adapter** and select **Ethernet Adapter**, **Fibre Channel Adapter**, **SCSI Adapter**, or **Serial Adapter**.

   c) In the **Virtual Adapters** page, click **Next**.

11. In the **SR-IOV Logical Ports** page, click **Actions** > **Create Logical Port** and select **Ethernet Logical Port**, or **RoCE Logical Port**. The **Add SR-IOV Logical Port** is displayed.

    a) From the **SR-IOV Port** list, select the SR-IOV logical port to create the corresponding logical port, and click **OK**.

    b) In the **SR-IOV Logical Ports** page, click **Next**.

12. In the **OptiConnect Settings** page, select the **Use virtual OptiConnect** check box to set virtual OptiConnect. Select **Use High Speed Link (HSL) OptiConnect** check box to set High Speed Link (HSL) OptiConnect.
13. In the **Tagged I/O** page, specify the load source that you want the logical partition to use as its load source from the **Load source** list. Select an alternate device on which the system restarts from the **Alternate restart device** list. Select the console from the **Console** list. Select an alternate console from the **Alternate console** list. Select the operations console from the **Operations Console** list, and click **Next**.
14. In the **Optional Settings** page, select the **Enable connection monitoring** check box to enable connection monitoring. Select the **Automatically start when the managed system is powered on** check box to start the partition profile automatically when the managed system is powered on. Select the **Enable redundant error path reporting** check box to receive reports about any redundant errors. Select **Enable electronic reporting of errors that cause partition termination or require attention** to receive electronic reports about any errors that cause logical partition termination or that require attention, and click **Next**.
15. The **Profile Summary** page displays the summary information about the logical partition and the partition profile. Click **Details** to view details about the physical I/O devices. Click **Finish** to create the logical partition and the partition profile.

Copying a partition profile

You can create a copy of an existing partition profile using the Hardware Management Console (HMC). After you create a copy of the existing partition profile, you can change the resource allocations within the new partition profile. This allows you to create multiple, nearly identical partition profiles without having to re-enter all of the resource allocations repeatedly.

About this task

To copy a partition profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to copy and click **Actions** > **Copy**.
5. Enter the name of the new partition profile into **New profile name** and click **OK**.

Changing partition profile properties

You can change the properties of a partition profile using the Hardware Management Console (HMC). Changing the properties of a partition profile changes the resource amounts that are assigned to a logical partition when you shut down and restart the logical partition using the changed partition profile.

Before you begin

A partition profile stores the required number of processors, memory, and hardware resources assigned to that profile. Any partition profile property changes are not applied to the logical partition until you activate the partition profile.

If you plan to change a partition profile that specifies dedicated memory to a partition profile that specifies shared memory, the following conditions apply:

- The HMC automatically deletes all of the physical I/O adapters specified in the partition profile. You can assign only virtual adapters to logical partitions that use shared memory.
- You must specify shared processors. Logical partitions that use shared memory must also use shared processors.

About this task

To change partition profile properties using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to change and click **Actions** > **Edit**.

   To add, remove, or change the vNIC adapter settings, you can run the **chsyscfg** command from the HMC command line. To add vNIC backing devices to a partition or to remove vNIC backing devices from a partition, and to change the vNIC auto-failback policy or to change the vNIC backing device failover policy, run the **chhwres** command from the HMC command line.

   When the HMC is at Version 9.1.0, or later, you can use the *max_capacity* field in the vNIC backing device attribute of the **chsyscfg** command to configure vNIC backing devices. You can also use the *max_capacity* attribute of the **chsyscfg** command to configure a single root I/O virtualization (SR-IOV) Ethernet logical port.
5. In the **General** tab, specify a name for the profile in the **Profile name** field, and click **OK**.
6. In the **Processors** tab, select **Shared** to assign partial processor units from the shared processor pool or select **Dedicated** to assign the entire processor that can only be used by the logical partition, and click **OK**.

   a) In the **Dedicated processors** section, specify values for the logical partition in the **Minimum dedicated processors**, **Desired dedicated processors**, and **Maximum dedicated processors** fields.

   b) In the **Processor Sharing** section, select **Allow when partition is inactive** check box to allow processor sharing when the partition is inactive. Select **Allow partition is active** check box to allow processor sharing when the partition is active. Select the processor compatibility mode from the **Processor compatibility mode** list, and click **OK**.
7. In the **Memory** tab, specify values for the logical partition memory in the **Minimum memory**, **Desired memory**, and **Maximum memory** fields, and click **OK**.

8. In the **I/O** tab, select the adapters from the adapter list to be included in the partition profile and click **Add as required**, or **Add as desired**. Click **Remove** to remove the selected adapter from the partition profile, and click **OK**.
9. In the **Virtual Adapters** tab, select the virtual adapter from the available adapter list and click **Actions** > **Properties** to view the properties of the virtual adapter.
   a) In the **Maximum virtual adapters** field, specify the number of maximum virtual adapters for the logical partition.
   b) To create a virtual adapter, click **Actions** > **Create Virtual Adapter** and select **Ethernet Adapter**, **Fibre Channel Adapter**, **SCSI Adapter**, or **Serial Adapter**.
   c) In the **Virtual Adapters** page, click **OK** to apply the changes.
10. In the **Power Controlling** tab, select the power controlling partitions from the **Power controlling partitions to add** list. Click **Add** to add the selected power controlling partition. Click **Remove** to remove the selected power controlling partition, and click **OK**.
11. In the **Settings** tab, select the **Enable connection monitoring** check box to enable connection monitoring. Select the **Automatically start when the managed system is powered on** check box to start the partition profile automatically when the managed system is powered on. Select the **Enable redundant error path reporting** check box to receive reports about any redundant errors. Select **Enable electronic reporting of errors that cause partition termination or require attention** to receive electronic reports about any errors that cause logical partition termination or that require attention.
    a) From the **Workload Management** section, select the partition workload group from the **Partition workload group** list, and click **OK**.
12. In the **SR-IOV Logical Ports** tab, click **SR-IOV Menu** > **Add Logical Port** and select **Ethernet Logical Port**, or **RoCE Logical Port**. The **Add SR-IOV Logical Port** page opens.
    a) From the **SR-IOV Port** list, select the SR-IOV logical port to create the corresponding logical port, and click **OK**.
    b) In the **SR-IOV Logical Ports** tab, click **OK**.
13. In the **SR-IOV Logical Ports** tab, a list of SR-IOV logical ports that are configured for the selected partition is displayed.
14. Select an SR-IOV logical port that you want to change, click **SR-IOV Menu** > **Edit Logical Port**. The **Logical Port Properties** page is displayed.
15. Select the **Diagnostic Mode** check box to enable or disable the setting.

    Note: Diagnostic mode can be set only if other logical ports are not associated with the physical port.
16. Select the **Promiscuous Mode** check box, if you want to enable the settings for the SR-IOV port. These settings are disabled by default.

    Note: You must select the **Promiscuous Mode** check box, if you want to virtualize the logical port even further, such as, if you want to use the logical port as the network adapter for shared Ethernet adapter (SEA).
17. Select the **Migratable** check box to configure SR-IOV logical ports in a client partition and to mark the SR-IOV logical ports as migratable by creating a new backup device, which can either be a virtual Ethernet adapter or a virtual NIC adapter.
    a) Select the **Configure a new back backup device** option to configure a new backup device. By default, this setting is enabled.
       i) Select **Virtual Ethernet Adapter** or **Virtual NIC Adapter** from the **Backup Device Type** option. By default, the **Virtual NIC Adapter** option is enabled.
       ii) Click **Configure backup device**. The **Configure SR-IOV Virtual NIC Migratable Backup** page is displayed.
          - Select a physical port that you want as a backup port for the migratable logical port from the list of available physical ports.
          - Select a hosting VIOS from the **Hosting Partition** list.

        - Select the appropriate capacity from the **Capacity** list.
        - Click **OK** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the window.

    iii) Optional: If you have selected **Virtual Ethernet Adapter** from the **Backup Device Type** option, the **Configure SR-IOV Virtual Ethernet Migratable Backup** page is displayed.

        - Select a virtual network that you want as a backup network for the migratable logical port from the list of available virtual networks.
        - Click **OK** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the window.

    Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

18. In the **Tagged I/O** tab, specify the load source that you want the logical partition to use as its load source from the **Load source** list. Select an alternate device on which the system restarts from the **Alternate restart device** list. Select the console from the **Console** list. Select an alternate console from the **Alternate console** list. Select the operations console from the **Operations Console** list, and click **OK**.
19. In the **OptiConnect tab**, select the **Use virtual OptiConnect** check box to set virtual OptiConnect. Select the **Use High Speed Link (HSL) OptiConnect** check box to set High Speed Link (HSL) OptiConnect, and click **OK**.

Deleting a partition profile

You can delete a partition profile using the HMC Hardware Management Console (HMC). This allows you to remove partition profiles that you no longer use.

Before you begin

Note: You cannot delete a partition profile that is the default partition profile for the logical partition. If the partition profile you want to delete is the default partition profile, you must first change the default profile to another partition profile.

About this task

To delete a partition profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions** > **Profiles** > **Manage Profiles**.
4. Select the partition profile that you want to delete and click **Actions** > **Delete**.
5. Click **OK** to confirm.

Managing virtual NICs on a logical partition

Learn how to manage virtual Network Interface Controllers (vNICs) on a partition.

You can use the Hardware Management Console (HMC) to complete the following tasks related to the virtual NICs on a partition:

- Adding virtual NICs
- Viewing virtual NICs
- Changing virtual NICs
- Removing virtual NICs

Adding virtual NICs

You can add virtual NICs to a partition by using the Hardware Management Console (HMC).

Before you begin

Before you add a virtual NIC, ensure that your system meets the following prerequisites, if the client partition is running:

- The Virtual I/O Server (VIOS) that hosts the virtual NIC is running with an active Resource Monitoring and Control (RMC) connection.
- The client partition has an active RMC connection.

Ensure that your system meets the following prerequisite, if the client partition is shut down:

- The Virtual I/O Server (VIOS) that hosts the virtual NIC is running with an active RMC connection or is shut down.

About this task

To add virtual NICs by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual NICs**. The **Virtual NIC** page opens.
5. Click **Add Virtual NIC**. The **Add Virtual NIC - Dedicated** page opens with the SR-IOV physical ports listed in a table.
6. Click **Add Entry** or **Remove Entry** to add or to remove backing devices for the Virtual NIC.

   Note: The **vNIC Auto Priority Failover** list is displayed when you add the second backing device entry. If you select **Enabled** from the **vNIC Auto Priority Failover** list, the hypervisor automatically fails over to the operational backing device that has the highest failover priority. Alternatively, if you select **Disabled**, the hypervisor does not take any action even if another operational backing device has a higher failover priority.
7. To configure each backing device entry, complete the following actions:
   a) Select the SR-IOV physical port on which you want to create the logical port to support the virtual NIC.

      Note: You must assign a different SR-IOV physical port for each backing device.

   b) Select the hosting partition.

   c) Specify the logical port minimum capacity.

      Note: The capacity of the logical port must be a percentage of the capacity of the SR-IOV physical port. If you do not specify a value, the HMC assigns the minimum capacity of the Ethernet logical port. The failover priority for the backing device must be in the range of 1 - 100, where 1 indicates the highest priority and 100 indicates the lowest priority. If you do not specify any value, the default priority value 50 is assigned to the backing device.

d) Specify the failover priority for the backing device.

8. Click **Advanced Virtual NIC Settings** to configure additional settings for the virtual NIC, such as virtual NIC adapter ID, MAC address settings, and the VLAN ID settings.
9. Click **OK**. The virtual NIC is added to the partition.

Viewing virtual NICs

You can view the properties of the virtual NIC backing device by using the Hardware Management Console (HMC).

About this task

To view the properties of the virtual NIC backing device by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual NICs**. The **Virtual NICs** page opens with the virtual NIC adapters listed in a table.
5. Select the virtual NIC from the list for which you want to view the properties.
6. Click **Action** > **View**. The **View Virtual NIC** page opens.
7. View the properties of the virtual NIC backing device, the MAC address settings, and VLAN ID settings for the virtual NIC.
8. Click **Close**.

Changing virtual NICs

You can change the properties of the virtual NIC by using the Hardware Management Console (HMC).

About this task

To change the properties of the virtual NIC by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual NICs**. The **Virtual NICs** page opens with the virtual NIC adapters listed in a table.
5. Select the virtual NIC from the list for which you want to change the properties.
6. Click **Action** > **Modify**. The **Modify Virtual NIC** page opens.

7. View the properties of the backing device, the MAC address settings, and VLAN ID settings for the virtual NIC.
8. You can change the Port VLAN ID and PVID priority for the selected virtual NIC.
9. Click **Close**.

Removing virtual NICs

You can remove the virtual NIC by using the Hardware Management Console (HMC).

About this task

To remove the virtual NIC by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual NICs**. The **Virtual NICs** page opens with the virtual NIC adapters listed in a table.
5. Select the virtual NIC that you want to remove.
6. Click **Action** > **Remove**. A delete confirmation message appears.
7. Click **OK** to remove the selected virtual NIC.

Managing virtual networks

Learn about managing PowerVM virtual networks on a partition.

You can use the Hardware Management Console (HMC) to complete the following networking tasks on a partition:

- Viewing virtual networks
- Changing virtual networks
- Removing virtual networks

Viewing the virtual network configuration

You can view the configuration details of the PowerVM virtual networks that are assigned to a partition by using the Hardware Management Console (HMC).

About this task

To view the configuration details of the PowerVM virtual networks by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.

3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Networks**. The **Virtual Networks** page opens.

Results

You can view the configuration details of the virtual networks in the table that is displayed in the **Virtual Networks** tab. The configuration details for each virtual network include the following information:

- Virtual Network Name
- VLAN ID
- Virtual Switch
- Network Bridge

Managing virtual network connections in adapter view

You can manage the PowerVM virtual network connections that are assigned to a partition in the adapter view, by using the Hardware Management Console (HMC).

About this task

To manage the virtual network connections in the adapter view by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Networks**. The **Virtual Networks** page opens in the **Network(s)** view. The currently available virtual network connections for the partition are listed in a table.
5. In the **Virtual Networks** work pane, click the right arrow key button to select **Adapter(s)** view. The currently available virtual Ethernet adapters and trunk adapters for the partition that are listed in a table. You can view, modify, or remove an adapter by using the **Action** menu.

   You can also create a trunk adapter for an IBM i logical partition.
6. To change the adapter settings, complete the following steps:

   a) Select the adapter for which you want to change the settings and click **Action** > **Modify Virtual Ethernet Adapter Settings**. Alternatively, you can right-click the adapter and select **Modify Virtual Ethernet Adapter Settings**. The **Modify Virtual Ethernet Adapter Settings** page opens with the virtual Ethernet adapter ID, virtual LAN (VLAN) ID, 802.1Q VLAN ID, and trunk priority of the adapter displayed in a table.

   b) Change the virtual Ethernet adapter settings including the Media Access Control (MAC) address settings, QoS settings, and 802.1Qbg settings, or IEEE settings and click **OK**.

   You can choose the following values for operating system-defined the MAC addresses:

   - **Allow all**: Allows any operating system-defined MAC addresses. This value is the default that is displayed.
   - **Deny all**: Does not allow any operating system-defined MAC addresses.

    - **Allow specified**: Specifies a maximum of four operating system-defined MAC addresses. You can add the MAC addresses to the **Allowed MAC Addresses** list.

    Notes:

    - Under **IEEE settings**, if you select the **IEEE 802.1q Compatibility** check box, additional VLANs can be supported on an Ethernet network. If you do not need additional VLANs for the trunk adapter, clear the **IEEE 802.1q Compatibility** check box.
    - The **802.11q VLAN ID** option is supported only for trunk adapters.

7. To view the adapter settings, complete the following steps:
    a) Select the adapter that you want to view and click **Action** > **View Virtual Ethernet Adapter Settings**. Alternatively, you can right-click the adapter and select **View Virtual Ethernet Adapter Settings**. The **View Virtual Ethernet Adapter Settings** page opens.
    b) View the virtual Ethernet adapter settings including the Media Access Control (MAC) address settings, QoS settings, and 802.1Qbg settings, or IEEE settings and click **Close**.
8. To remove the adapter settings, complete the following steps:
    a) Select the adapter that you want to remove and click **Action** > **Remove Virtual Ethernet Adapter Settings**. Alternatively, you can right-click the adapter and select **Remove Virtual Ethernet Adapter Settings**.
    b) When you are prompted to confirm the removal, click **OK**.

Creating trunk adapters

When the Hardware Management Console (HMC is at version 8.7.0, or later, you can create trunk adapters for an IBM i logical partition.

About this task

To create a trunk adapter by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Networks**. The **Virtual Networks** page opens. You can use the left and right arrow key buttons to switch between **Network(s)** and **Adapter(s)** views. The **Network(s)** view is the default view.
5. In the **Virtual Networks** work pane, click the right arrow key button to select **Adapter(s)** view. The currently available virtual Ethernet adapters and trunk adapters for the partition are listed in a table.
6. Click **Create Trunk Adapter**. The **Create Trunk Adapter** page is displayed. You can add a trunk adapter to an IBM i logical partition.
7. Optional: Click **View Existing Virtual Networks** to view the list of all existing virtual networks on the managed system.
8. In the **Port VLAN ID** field, enter the VLAN ID on which the trunk adapter must be created.
9. In the **Virtual Switch** field, select a virtual switch that can be assigned to the trunk adapter from the list of virtual switches that are configured on the managed system.
10. In the **Trunk Priority** field, set the trunk priority for the trunk adapters to either 1 or 2.
11. In the **MAC Address** field, enter the MAC address.
12. In the **OS MAC Address Restrictions** field, specify the MAC address restrictions. The available values are **Allow all**, **Deny all**, and **Allow Specified**.

13. In the **QoS Settings** section, select the **Enable QoS Setting** option, if you want to provide different priorities to different applications, users, or data flows to maintain the performance of the network. The value for the priority level ranges in the range 0 - 7.
14. In the **IEEE Settings** section, select **IEEE 802.11q Compatibility** if you need additional VLANs to be supported on an Ethernet network. Alternatively, you can clear the **IEEE 802.1q Compatibility** check box if you do not need additional VLANs for the trunk adapter.

    If you select **IEEE 802.1q Compatibility**, the **802.11q VLAN ID** field is displayed. You can select additional VLAN IDs.
15. Click **OK**.

Managing virtual network connections in network view

You can manage the PowerVM virtual network connections that are assigned to a partition in the network view by using the Hardware Management Console (HMC).

About this task

To manage the virtual network connections in the network view by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Networks**. The **Virtual Networks** page is displayed. You can use the left and right arrow key buttons to switch between **Network(s)** and **Adapter(s)** views. The **Network(s)** view is the default view. The currently available virtual network connections for the partition are listed in a table, with the details about the virtual network name, VLAN ID, virtual switch, virtual network bridge, and virtual Ethernet adapter ID that are associated with each virtual network. You can attach new virtual networks by clicking **Attach Virtual Network**. You can also detach an existing virtual network by selecting the virtual network that you want to remove and by using the **Action** menu.
5. To attach a virtual network click **Attach Virtual Network**. All virtual networks that are discovered by the managed system are displayed in a table. The table lists all the assigned virtual networks, the VLAN ID, the virtual switch, the network bridge name to which the virtual network is attached.
6. To attach an existing virtual network to a partition, complete the following steps:
   a) Select **Show and attach new virtual Ethernet adapters**.
   b) Select the check box near the name of the virtual network that you want to connect to the logical partition.
   c) In the **Virtual Ethernet Adapter ID** field, enter the virtual Ethernet adapter ID.
   d) Click **OK**.
7. To remove a virtual network from the partition, complete the following steps:
   a) In the **Virtual Networks** table, select the virtual network that you want to remove and click **Action** > **Detach**. Alternatively, you can right-click the virtual network and select **Detach**. The **Detach Virtual Network** page is displayed.
   b) In the **Virtual Ethernet Adapters** table, select the check box near the name of the virtual network that you want to remove.
   c) Click **OK**.

Managing virtual storage for a partition

Learn about managing virtual storage for a partition.

You can use the Hardware Management Console (HMC) to complete the following storage tasks on a partition:

- Managing virtual storage for a partition in Adapter view.
  - Creating virtual SCSI and virtual Fibre Channel adapters
  - Removing virtual SCSI and virtual Fibre Channel adapters
- Managing virtual storage for a partition in Storage view.
  - Managing virtual SCSI resources for a partition
  - Viewing virtual Fibre Channel assignment
  - Optical device assignment

Managing virtual storage for a partition in adapter view

You can create, view, and manage virtual storage allocated to a partition by using the Hardware Management Console (HMC).

You can add the required virtual storage resources to a partition. In the **Virtual Storage** work pane, you can use the left and right arrow key buttons to switch between **Storage View** and **Adapter View**. Click the right arrow key button to select **Adapter View**. In the **Adapter View**, you can view the adapter configuration of the virtual storage devices that are allocated for the logical partition. The **Adapter View** provides a mapping of the adapters to the physical storage device in a logical partition.

In the **Adapter View**, you can create, view, and manage the properties of virtual Small Computer Serial Interface (SCSI) and Virtual Fibre Channel (VFC) adapters for the partition on the managed system. You can also add different types of storage devices to the logical partition. To launch the **Add Virtual SCSI Device** page, select an adapter from the list available in the table and click **Action** > **Add Client Device**. Alternatively, you can right-click the adapter and select **Add Client Device**. For more information, see "Managing virtual SCSI resources for a partition" on page 93.

Creating IBM i or Virtual I/O Server hosted Virtual SCSI adapters

When the Hardware Management Console (HMC is at version 8.7.0, or later, you can view and manage the IBM i or Virtual I/O Server hosted virtual SCSI adapters for the partition.

To add an IBM i or Virtual I/O Server hosted virtual SCSI adapter, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
5. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**.
6. In the **Virtual Storage** work pane, click the right arrow key button to select **Adapter View**. The **Virtual SCSI Adapters** tab is displayed by default.
7. In the **IBM i / Virtual I/O Server Virtual SCSI Adapters** section, click **Create Adapter**. The **Create Virtual SCSI Adapter** window is displayed.
8. From the **Remote Partition Type** option, select **IBM i** or **Virtual I/O Server**.
9. In the **Adapter ID** field, enter the adapter ID.

Note: If you do not want to specify an adapter ID, you can continue the procedure with the adapter ID that is populated automatically in the **Adapter ID** field. The adapter ID displayed in this field is the next available slot ID for the virtual SCSI client adapter that is being created.

10. From the **Remote Partition** list, select an IBM i or Virtual I/O Server partition to which the virtual SCSI client adapter connects.
    - If you have selected IBM i from the **Remote Partition Type** option, select an IBM i partition from the **Remote Partition** list. The list displays all the IBM i partitions that are available in the managed system for creating the virtual SCSI adapter.
    - If you have selected **Virtual I/O Server** from the **Remote Partition Type** option, select a virtual I/O server partition from the **Remote Partition** list. The list displays all the VIOS partitions that are available in the managed system for creating the virtual SCSI adapter.
11. By default, the **Remote Adapter** check box is enabled which means both the client and the server adapters are created. If you want to create only the client adapter, clear the **Remote Adapter** check box.
12. If the client partition is an IBM i partition, and you want to create only the server adapter, you can select the adapter type. From the **Adapter Type** option, select **Client** or **Server**. The **Remote Partition Type** is disabled if you select **Server** from the **Adapter Type** option.
13. From the **Remote Adapter ID** list, select the remote adapter ID. The remote slot number of the selected IBM i or Virtual I/O Server partition is displayed in the **Remote Partition ID** field. This field is populated automatically with the next available slot ID, which is based on the IBM i or Virtual I/O Server partition that is selected for creating the virtual SCSI server adapter. Alternatively, you can click **Populate existing usable Remote Adapter IDs**. All server adapters that exist in the selected IBM i or Virtual I/O Server partition, and which are not connected to any logical partition, are displayed in the **Remote Adapter ID** field.
14. Click **OK** to apply the changes. Alternatively, you can click **Cancel** to reject the changes and to close the window.

You can remove the virtual SCSI adapter that is configured for a logical partition. Select the virtual SCSI adapter that you want to remove from the **Virtual SCSI Adapters** list, and click **Action** > **Remove**. Click **OK** when prompted.

Creating Virtual I/O Server hosted Virtual Fibre Channel Adapter

To create a Virtual I/O Server hosted Virtual Fibre Channel (VFC) adapter, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
5. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**.
6. In the **Virtual Storage** work pane, click the right arrow key button to select **Adapter View**. The **Virtual SCSI Adapters** tab is displayed by default.
7. Select **Virtual Fibre Channel Adapters** tab. A list of virtual Fibre Channel assignments to a partition is displayed.
8. In the **Virtual FC Adapters** section, click **Create Adapter**. The **Create Virtual Fibre Channel Adapter** window is displayed.
9. In the **Adapter ID** field, enter the adapter ID.

Note: If you do not want to specify an adapter ID, you can continue the procedure with the adapter ID that is populated automatically in the **Adapter ID** field. The adapter ID displayed in this field is the next available slot ID for the virtual Fibre Channel adapter that is being created.

10. From the **Remote Partition** list, select a Virtual I/O Server partition to which the virtual Fibre Channel adapter connects. The list displays all the VIOS partitions that are available in the managed system for creating the virtual Fibre Channel adapter.
11. From the **Remote Adapter ID** list, select the remote adapter ID. The remote slot number of the selected Virtual I/O Server partition is displayed in the **Remote Partition ID** field. This field is populated automatically with the next available slot ID, which is based on the Virtual I/O Server partition that is selected for creating the virtual Fibre Channel adapter. Alternatively, you can click **Populate existing usable Remote Adapter IDs**. All server adapters that exist in the selected Virtual I/O Server partition, and which are not connected to any logical partition, are displayed in the **Remote Adapter ID** field.
12. Click **OK** to apply the changes. Alternatively, you can click **Cancel** to reject the changes and to close the window.

You can remove the virtual Fibre Channel adapter that is configured for a logical partition. Select the virtual Fibre Channel adapter that you want to remove from the **Virtual FC Adapters** list, and click **Action** > **Remove**. Click **OK** when prompted.

Managing virtual storage for a partition in storage view

You can create, view, and manage virtual storage allocated to a partition by using the Hardware Management Console (HMC).

You can add the required virtual storage resources to a partition. In the **Virtual Storage** work pane, you can use the left and right arrow key buttons to switch between the **Storage View** and **Adapter View**. Click the left arrow key button to select **Storage View**. In the **Storage View**, you can view and manage the storage capability of the logical partition. The **Storage View** is the default view.

In the **Storage View**, you can view and manage the virtual SCSI adapters, VFC adapters, and the virtual optical devices that are configured to a logical partition.

Managing virtual SCSI resources for a partition

You can assign virtual Small Computer Serial Interface (SCSI) resources to a partition by using the Hardware Management Console (HMC).

Using the virtual SCSI adapter, client partitions can share disk storage and optical devices that are assigned to the Virtual I/O Server (VIOS) partition.

You can add different types of virtual storage to the PowerVM® configuration such as, **Physical Volume**, **Shared Storage Pool Volume**, or **Logical Volume**. By default, the **Physical Volume** table is displayed.

You can view the device mapping details of the storage devices in a logical partition. Right-click on the storage device and select View Device Mapping. The storage device details and the connected Virtual I/O Server details are displayed.

You can also add the Virtual I/O Server to provide an adapter connection. Click Edit Connections and select the Virtual I/O Server and the server adapters to provide the adapter connection.

The Virtual SCSI tab displays the end to end mapping for the virtual SCSI that includes the server adapter, client adapter, and the storage that is used by the virtual SCSI adapter that is configured for a particular logical partition. You can also remove the client or server adapter that is configured for the particular partition.

Adding virtual SCSI devices

You can add different types of virtual storage, such as a physical volume, a shared storage pool volume, or a logical volume from the Virtual SCSI tab in the Storage View. Only the virtual storage devices assigned to the PowerVM configuration are displayed here.

Adding a Physical Volume

To add a physical volume, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
5. In the **Properties** page, click **Virtual I/O** > **Virtual Storage**.
6. In the **Virtual SCSI** tab, click **Add Physical Volume**. The **Add Physical Volume** page opens and the physical volumes table is displayed.
7. Select physical volumes from the list available in the table.

   Note: You can select the **Show in use physical volumes** check box to see the assigned physical volumes. You can also click **Run ConfigDevice** to refresh the list of physical volumes
8. Specify the Target Name for the physical volume that you want to add to the partition.
9. Click **Edit Connection** if you want to change the server Adapter ID and client Adapter ID that are assigned to the physical volume. The **Edit Connection** page opens. You can enter a server Adapter ID and client Adapter ID.
10. Select the **Virtual I/O servers** to provide an adapter connection. You can select up to three Virtual I/O Servers.
11. Enter a **Server Adapter ID** and a **Client Adapter ID** that you want to assign for the adapter connection, or click **Use Existing Adapters** to select a server adapter ID from the list.

    Note: If you clicked **Use Existing Adapters**, you cannot modify the client Adapter ID as the client Adapter ID is automatically assigned to the VIOS.
12. Click **OK** to assign the physical volume to the partition.

Adding a Shared Storage Pool Volume

To add a Shared Storage Pool volume, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
5. In the **Properties** page, click **Virtual I/O** > **Virtual Storage**.
6. In the **Virtual SCSI** tab, click **Shared Storage Pool Volume**. The **Shared Storage Pool Volumes** table is displayed.
7. Click **Add Shared Storage Pool Volume**. The **Add Shared Storage Pool Volume** page is displayed.
8. Select a **Storage Cluster** from the list.
9. Select **Add new Shared Storage Pool Volume** to add a Shared Storage Pool volume or select **Add existing Shared Storage Pool volume**.

- If you chose to add a new Shared Storage Pool volume, select the tier to which the new SSP volume is associated and enter a device name and size. Select the VIOS connections that are assigned to the PowerVM configuration.
- If you chose to add an existing Shared Storage Pool (SSP) volume, select an existing SSP volume. Select the VIOS connections that are assigned to the PowerVM configuration.

10. Click **Edit Connection** if you want to change the server Adapter ID and client Adapter ID that are assigned to the shared storage pool volume. The **Edit Connection** page opens. You can enter a server Adapter ID and client Adapter ID.
11. Select the **Virtual I/O servers** to provide an adapter connection. You can select up to three Virtual I/O Servers.
12. Enter a **Server Adapter ID** and a **Client Adapter ID** that you want to assign for the adapter connection, or click **Use Existing Adapters** to select a server adapter ID from the list.

    Note: If you clicked **Use Existing Adapters**, you cannot modify the client Adapter ID as the client Adapter ID is automatically assigned to the VIOS.
13. Select **Shared Storage Pool Volumes** to view the SSP volumes that are already assigned to the existing logical partition
14. Click **OK** to assign the Shared Storage Pool volume to the partition.

Adding a Logical Volume

To add a Logical volume, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select a system and click **Actions** > **View System Partitions**. In the **Partitions** page, you can view all the partitions that belong to the system.
4. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
5. In the **Properties** page, click **Virtual I/O** > **Virtual Storage**.
6. In the **Virtual SCSI** tab, click **Logical Volume**. The **Logical Volume** table is displayed.
7. Click **Add Logical Volume** to add logical volumes to a partition. The **Add Logical Volume** page is displayed.
8. Select a volume group from the table.
9. Select **Add new logical volume** to add a logical volume or select **Add an existing logical volume**.
    - If you chose to add a new logical volume, enter a device name and size.
    - If you chose to add an existing logical volume, select an existing device name.
10. Click **Edit Connection** if you want to change the server Adapter ID and client Adapter ID that are assigned to the logical volume. The **Edit Connection** page opens. You can enter a server Adapter ID and client Adapter ID.
11. Select the **Virtual I/O servers** to provide an adapter connection. You can select up to three Virtual I/O Servers.
12. Enter a **Server Adapter ID** and a **Client Adapter ID** that you want to assign for the adapter connection, or click **Use Existing Adapters** to select a server adapter ID from the list.

    Note: If you clicked **Use Existing Adapters**, you cannot modify the client Adapter ID as the client Adapter ID is automatically assigned to the VIOS.
13. Click **OK** to assign the Logical volume to the partition.

Viewing virtual Fibre Channel assignments to a partition
You can view the virtual Fibre Channel resources that are assigned to a partition by using the Hardware Management Console (HMC).

About this task

To view the virtual Fibre Channel resources that are assigned to a partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page is displayed.
5. Click the **Virtual Fibre Channel** tab. The **Virtual Fibre Channel** page opens. You can use the left and right arrow key buttons to switch between the **Storage View** and **Adapter View**. The **Storage View** view is the default view. The table displays all the virtual Fibre Channel resources assigned to the partition. Only the virtual Fibre Channel resources assigned to the PowerVM configuration are displayed.
6. In the **Virtual Storage** work pane, click the right arrow key button to select **Adapter View**.
7. Select **Virtual Fibre Channel Adapters** tab. A list of virtual Fibre Channel assignments to a partition is displayed.

Assigning virtual Fibre Channel storage to a partition
You can assign virtual Fibre Channel storage to a partition by using the Hardware Management Console (HMC).

About this task

To assign virtual Fibre Channel storage to a partition by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Virtual Fibre Channel** tab. The **Virtual Fibre Channel** page opens in **Storage View**.
6. Click **Add Virtual Storage**. The **Add Virtual Fibre Channel** page opens.
7. Select a Virtual storage type from the **Virtual I/O Server** list.
8. Select Fibre Channel ports from the **Fibre Channel port** list available in the selected **VIOS**.

   Note: You can click **Edit connection** to manually configure the **Virtual Fibre Channel** adapter settings for the connection. Enter the **WWPN** details and the **Server Adapter ID**.
9. Click **Save**. The Fibre Channel port is assigned to the partition.

Optical device assignment

You can manage the optical devices that are assigned to partitions by using the Hardware Management Console (HMC).

Viewing physical optical devices

You can view the physical optical devices that can be assigned to a partition by using the Hardware Management Console (HMC).

About this task

To view the physical optical devices by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Physical Optical Device** tab. A list of physical optical devices that can be assigned to the selected partition is displayed.
6. To view the mapping of a physical optical device, click on a physical optical device and select **View Device Mapping**.

   In the **Physical Optical Device** area, you can view details like the device name, description, and the physical location. In the **Virtual I/O Server** area, you can view the client adapter name, and the server adapter name.
7. Click **Close**.

Adding physical optical devices

You can add physical optical devices to a partition by using the Hardware Management Console (HMC).

About this task

To add physical optical devices by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Physical Optical Device** tab.
6. Click the **Add Virtual Storage** tab. The **Add Physical Optical Device** page opens.
7. Select the physical optical device to be added to the PowerVM configuration. Click **OK**.

   Only devices that are assigned through the PowerVM configuration are displayed and you can select a device only from the displayed list.

8. To edit the properties of the device, click **Edit Connections**. You can select up to three Virtual I/O Servers to provide an adapter connection.
9. For each Virtual I/O Server, select a value from the **Server Adapter ID** list.
10. Click **OK**. The physical optical device is added to the partition.
11. Click **Close**.

Removing physical optical devices

You can remove a physical optical devices that is assigned to a partition by using the Hardware Management Console (HMC).

About this task

To remove a physical optical device by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Physical Optical Device** tab.
6. Select a physical device and click **Remove**.
7. When the device to be removed is assigned to a running partition, you are prompted with a message verifying whether you want to continue to remove the device.
8. Click **OK** to remove the device, or click **Cancel** to exit the operation.
9. Click **Close**.

Viewing virtual optical devices

You can view the virtual optical devices that can be assigned to a partition by using the Hardware Management Console (HMC).

About this task

To view the virtual optical devices by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Virtual Optical Device** tab. A list of virtual optical devices that are assigned to the selected partition is displayed.
6. To view the mapping of a virtual optical device, click a virtual optical device and select **View Device Mapping**.

In the **Virtual Optical Device** area, you can view details such as the device name, the media file, and the size in GB. In the **Virtual I/O Server** area, you can view the client adapter name and the server adapter name.

7. Click **Close**.

Adding virtual optical devices

You can add virtual optical devices to a partition by using the Hardware Management Console (HMC).

About this task

To add a virtual optical device, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Virtual Optical Device** tab.
6. Click the **Add Virtual Storage** tab. The **Add Virtual Storage** page opens.
7. In the **Device Name** field, enter the device name and select the Virtual I/O Server from the table.
8. Click **OK**.
9. Optional:
10. You can select the server adapter ID to provide an adapter connection. Otherwise, the next available server adapter ID is used.
    a) To select the server adapter ID, click **Edit Connections**.
    b) From the **Server Adapter ID** list, select the server adapter ID.
11. Click **OK**. The virtual optical device is added to the partition.
12. Click **Close**.

Removing virtual optical devices

You can remove a virtual optical device that is assigned to a partition by using the Hardware Management Console (HMC).

About this task

To remove a virtual optical device by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Virtual Optical Device** tab.

6. Select a virtual device and click **Remove** When the device to be removed is assigned to a running partition, you are prompted with a message to verify whether you want to continue to remove the device.
7. Click **OK** to remove the device, or click **Cancel** to exit the operation.
8. Click **Close**.

Loading and unloading media files

You can load or unload media files to or from virtual optical devices by using the Hardware Management Console (HMC).

About this task

To load or unload a media file to or from a virtual optical device by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Virtual Storage**. The **Virtual Storage** page opens.
5. Click the **Virtual Optical Device** tab.
6. Select a virtual device and click **Load**.
7. Select the media file to assign to partition and click **OK**.

   Note: If there is a mount error, a message is displayed.
8. Click **Close**.
9. To remove a media file that is assigned to a partition, select the virtual optical device and click **Unload**.

Managing hardware virtualized I/O adapters

You can view and change the settings of hardware virtualized I/O adapters, such as single root I/O virtualization (SR-IOV) port adapters and logical host Ethernet adapters (LHEA) for a partition by using the Hardware Management Console (HMC).

SR-IOV logical port settings

You can add, change, and remove single root I/O virtualization (SR-IOV) logical ports that are configured on a partition by using the Hardware Management Console (HMC).

Adding SR-IOV logical ports

You can add single root I/O virtualization (SR-IOV) logical ports to a partition by using the Hardware Management Console (HMC).

Before you begin

About this task

To add an SR-IOV port to a partition by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. On the **SR-IOV** tab, click **Add Port**. The **Add SR-IOV Logical Port** page opens.
6. In the **Add SR-IOV Logical Port** page, select **Ethernet** or **RoCE** from the **Select Logical Port Type** option button.
7. Click **Select an SR-IOV physical port**. The **Physical Ports** page is displayed.
8. In the **Physical Ports** page, select a physical port from the list of available physical ports, and click **OK**.
9. In the **Logical port capacity** field, enter the capacity value in percentage for the logical port.

   Note: The sum of percentage of capacity values for all the configured logical ports on a physical port must be less than or equal to 100%. To minimize the configuration effort when you add more logical ports, you can reserve some capacity for the additional logical ports.
10. Expand **Advanced settings** to view the advanced setting options for the SR-IOV adapter.
11. Select the **Promiscuous Mode** check box, if you want to enable the settings for the SR-IOV port. These settings are disabled by default.

    Note: You must select the **Promiscuous Mode** checkbox, if you want to virtualize the logical port even further, such as, if you want to use the logical port as the network adapter for shared Ethernet adapter (SEA).
12. Select the **Migratable** check box to configure SR-IOV logical ports in a client partition and to mark the SR-IOV logical port as migratable by creating a new backup device, which can either be a virtual Ethernet adapter or a virtual NIC adapter.

    a) Select the **Configure a new back backup device** option to configure a new backup device. By default, this setting is enabled.

       i) Select **Virtual Ethernet Adapter** or **Virtual NIC Adapter** from the **Backup Device Type** option. By default, the **Virtual NIC Adapter** option is enabled.

       ii) Click **Configure backup device**. The **Configure Virtual NIC Backing Device** page is displayed.

       - In the **Physical Port Location Code** tab, select a physical port that you want as a backup port for the migratable logical port from the list.
       - In the **Hosting Partition** tab, select a hosting VIOS from the hosting partition list.
       - In the **Capacity** tab, select the appropriate capacity from the capacity list.
       - Click **OK** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the window.

       iii) Optional: If you have selected **Virtual Ethernet Adapter** from the **Backup Device Type** option, the **Attach Virtual Network** page is displayed.

       - Select a virtual network that you want as a backup network for the migratable logical port from the list of available virtual networks.
       - Click **OK** to apply the changes. Alternatively, click **Cancel** to reject the changes and to close the window.

    Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version

9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

13. From the **OS MAC Address Restrictions** list, select an option for the OS MAC address restrictions.
14. From the **VLAN ID Restrictions** list, select an option for the OS VLAN ID restrictions.
15. In the **Port VLAN ID** field, enter a value. The valid range is 2 - 4094.

    Note: The default value of Port VLAN ID is 0. If you enter a non-zero value in the Port VLAN ID field, the 802.1Q Priority field becomes available.
16. In the **802.1Q Priority** field, enter any value from 0 - 7, where 0 indicates the lowest priority and 7 indicates the highest priority value.
17. Click **OK**. The SR-IOV port is added to the partition.

Changing SR-IOV logical ports

You can change the settings of single root I/O virtualization (SR-IOV) logical ports on a partition by using the Hardware Management Console (HMC).

About this task

To change the settings of an SR-IOV port by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. Click the **SR-IOV** tab. A list of SR-IOV logical ports that are configured for the selected partition is displayed.
6. Right click an SR-IOV logical port that you want to change and select **Modify Logical Port**. The **Modify SR-IOV Logical Port** page opens.

   Note: Diagnostic mode can be set only if other logical ports are not associated with the physical port.
7. Select the **Diagnostic Mode** check box to enable or disable the setting.
8. If the **OS MAC Address Restrictions** option indicates **Allow Specified**, you can add MAC addresses to the **Specify allowed MAC Addresses** list.
9. If the **VLAN ID Restrictions** option indicates **Allow Specified**, you can add VLAN IDs to the **Specified VLAN IDs or range** list.
10. In the **Port VLAN ID** field, enter a value to change the existing value. The valid range is 2 - 4094.

    Note: The default value of Port VLAN ID is 0. If you enter a non-zero value in the Port VLAN ID field, the 802.1Q Priority field becomes available.
11. In the **802.1Q Priority** field, enter any value from 0 - 7, where 0 indicates the lowest priority and 7 indicates the highest priority value.
12. Click **OK** to save the changes that you made for the SR-IOV logical port settings.

Removing SR-IOV logical ports

You can remove single root I/O virtualization (SR-IOV) logical ports from a partition by using the Hardware Management Console (HMC).

About this task

To remove an SR-IOV port by using an HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. Click the **SR-IOV** tab. A list of SR-IOV logical ports that are configured for the selected partition is displayed.
6. Right-click an SR-IOV logical port that you want to remove and select **Remove Logical Port** > **OK**.

   Note: If the selected partition is powered on, the SR-IOV logical port needs to be deconfigured in the selected partition before being removed.

Results

The selected SR-IOV logical port is removed.

Logical host Ethernet adapter (LHEA) settings

You can view, add, change, and remove logical host Ethernet adapters (LHEAs) that are configured on a partition by using the Hardware Management Console (HMC).

An LHEA is a representation of a physical HEA on a partition. An LHEA appears to the operating system as a physical Ethernet adapter, just as a virtual Ethernet adapter appears as a physical Ethernet adapter. Each partition can have one LHEA for each physical HEA on the managed system. Each LHEA can have one or more logical ports, and each logical port can connect to a physical port on the HEA.

Adding logical host Ethernet adapters

You can add logical host Ethernet adapters (LHEAs) to a partition by using the Hardware Management Console (HMC).

About this task

You can select an LHEA from the list and add it to the partition with the required settings.

Procedure

To add an LHEA adapter to a partition, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.

4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. Click the **HEA** tab.
6. Click **Add Adapter**. The **Add LHEA Adapter** page opens.
7. From the list of physical ports, select the physical port to associate with the LHEA adapter. The list of physical ports is not displayed if there are no available ports.
8. Expand **Advanced Settings**.
9. From **MAC Address settings** option, select the MAC address settings.
10. From **VLAN ID Settings** option, select the VLAN ID settings .

    Note: The advanced settings are available only if the partition is QoS capable.
11. Click **OK**.

Results

The LHEA adapter is added to the partition.

Modifying logical host Ethernet adapter ports

You can change the settings of logical host Ethernet adapter (LHEA) ports on a partition by using the Hardware Management Console (HMC).

About this task

You can select an LHEA from the list and change it with the required settings.

Procedure

To change LHEA port settings, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. Click the **HEA** tab. A list of LHEAs configured for the selected partition is displayed.
6. Right-click an LHEA adapter that you want to change and select **Modify Port**. The **Modify Logical Host Ethernet Adapter Port** page opens.
7. Select the **Dedicated Mode** check box if you want to make the LHEA port dedicated to the assigned partition.
8. From the **MAC Address Settings** list, change the MAC address settings.
9. From the **VLAN ID Settings** list, change the VLAN ID settings.
10. Click **OK** to save the changes for the LHEA port.

Results

The LHEA port settings are saved.

Removing logical host Ethernet adapter ports

You can remove logical host Ethernet adapter (LHEA) ports from a partition by using the Hardware Management Console (HMC).

About this task

You can select an LHEA from the list and remove it from the partition.

Procedure

To remove an LHEA port, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. Click the **HEA** tab. A list of LHEA ports that are configured for the selected partition is displayed.
6. Right-click an LHEA port and select **Remove Port**.
7. Click **OK**. The selected LHEA port is removed after confirmation.

Results

The selected LHEA port is removed.

Managing host channel adapters on a partition

Host channel adapters (HCAs) provide port connections from a managed system to other devices. You can connect the port to another HCA, a partition, or a switch that redirects the incoming data from one port to a device that is attached to another port.

About this task

You can view a list of the HCAs on a partition that is managed by the Hardware Management Console (HMC). You can select an HCA from the list to display the current partition usage for the HCA.

Procedure

To manage HCA settings, complete the following steps:

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the partition for which you want to view or change the properties and capabilities and click **Actions** > **View Partition Properties**. The **Properties** page is displayed. You can view and change the properties that are listed under the **Properties** area.
4. In the **Properties** pane, click **Virtual I/O** > **Hardware Virtualized I/O**. The **Hardware Virtualized I/O** page opens.
5. In the work pane, click the **HCA** tab.
6. Click **Launch Manage Host Channel Adapters**. A window opens with the list of HCAs in a table.
7. From the table, select an HCA to display the current partition usage.
8. Click **OK**.

Viewing topology diagrams of a system

Learn how to view all the topology diagrams of a system.

You can use the Hardware Management Console (HMC) to view the topology diagrams of a system.

Viewing virtual networking diagrams

You can view the end-to-end network configuration for the selected system, by using the HMC. The view of the virtual networks begins with the physical adapter cards and the physical ports that are connected to them. As you scroll down, you can see the defined virtual bridges, link aggregation devices, virtual switches, virtual networks, and partitions in the VIOS.

About this task

You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the network diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the end-to-end network configuration for the selected system by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**.

   The **All Systems** page is displayed.
3. In the work pane, select the system in which the partition is located and click **Actions** > **View System Partitions**. The configuration page opens. You can view the configuration details of the system you selected.
4. In the navigation pane, click **Topology** > **Virtual Networking Diagram** to view the end-to-end network configuration for the selected system.
5. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
6. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
7. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols used in the virtual networking diagram.

Viewing virtual storage diagrams

Two types of virtual storage diagrams are available - systems storage and partition storage. You can view the virtual storage configuration for the selected system, including the physical and virtual components of system storage, by using the HMC. You can also view the virtual storage configuration for a single partition in a particular system, including the physical and virtual components of storage assigned to that particular partition, by using the HMC.

About this task

This diagram displays a high-level overview of the contents of the system or a single partition, and not specific component relationships. You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the storage diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the virtual storage configuration for the selected system or a single partition by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**.

   The **All Systems** page is displayed.
3. In the work pane, select the system in which the partition is located and click **Actions** > **View System Partitions**. The configuration page opens. You can view the configuration details of the system you selected.
4. In the navigation pane, click **Topology** > **Virtual Storage Diagram** to view the virtual storage configuration for the selected system.

   Note: To view the virtual storage diagrams of a single partition in a particular system, select the partition of your choice and then click **Topology** > **Partition Virtual Storage Diagram**.
5. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
6. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
7. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols used in the virtual storage diagram.

Viewing SR-IOV and vNIC diagrams

You can view the SR-IOV and virtual Network Interface Controllers (vNIC) configuration for the selected system, including the physical and virtual components, by using the HMC.

About this task

This diagram displays the relationships between SR-IOV adapters and other virtual components, such as vNIC. You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the SR-IOV and vNIC diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the SR-IOV and vNIC configuration for the selected system by using the HMC, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**.

   The **All Systems** page is displayed.
3. In the work pane, select the system in which the partition is located and click **Actions** > **View System Partitions**. The configuration page opens. You can view the configuration details of the system you selected.
4. In the navigation pane, click **Topology** > **SR-IOV and vNIC Diagram** to view the SR-IOV and vNIC configuration for the selected system.
5. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
6. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
7. In the upper-right corner of the work pane, click the **legend** icon to view an explanation of the symbols used in the SR-IOV and vNIC diagram.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled “Cookies, Web Beacons and Other Technologies”.

Programming interface information

This Managing the virtualization environment publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.2, IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.4, and IBM Virtual I/O Server Version 3.1.2.

Trademarks

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Monitoring the virtualization environment

Note

Before using this information and the product it supports, read the information in “Notices” on page 19.

This edition applies to IBM® AIX® Version 7.2, to IBM AIX Version 7.1, to IBM AIX Version 6.1, to IBM i 7.4 (product number 5770-SS1), to IBM Virtual I/O Server Version 3.1.1, and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

Contents

Monitoring the virtualization environment by using the Performance and Capacity Monitor function

The Performance and Capacity Monitor function collects allocation and usage data for virtualized server resources. When the Hardware Management Console is at version 8.6.0, or later, you can also export data metrics that are collected for the specified time. It displays data in the form of graphs and tables, which are viewable from the Performance and Capacity Monitor home page. The Performance and Capacity Monitor function is available in the Hardware Management Console (HMC) Version 8, Release 8.1.0, or later.

The Performance and Capacity Monitor function gathers and displays capacity reporting data and performance monitoring data. You can monitor processor, memory, virtual storage, and virtual network resource usage. This data can help you better understand how managed systems and logical partitions are using resources, and whether resources are under-used or over-used. It can also help you identify and fix performance bottlenecks. By using the Performance and Capacity Monitor, you can manage current capacity and plan for future requirements.

What's new in Monitoring the virtualization environment

Read about new or significantly changed information in Monitoring the virtualization environment since the previous update of this topic collection.

October 2019

- The following topics were updated with information about the support for the persistent memory:
  - "Current Resource Utilization graphs" on page 4
  - "The Resource Utilization table" on page 7
  - "Memory breakdown table" on page 10

August 2018

The following updates have been made to the content:

- Added information about PCM status and adapter statistics in the "Troubleshooting the Performance and Capacity Monitor" on page 14 topic.
- Removed or updated obsolete information in various topics.

Getting started

Learn how to use the Performance and Capacity Monitor.

To use the Performance and Capacity Monitor, refer to the following topics.

Enabling data collection

Server resource utilization monitoring starts after you enable data collection and continues until you disable it. Server utilization data is stored on the Hardware Management Console (HMC).

About this task

The number of servers that can be monitored is determined by the HMC. You can view the number of managed servers that can be monitored by accessing the Performance and Capacity Monitor preferences REST API at the following uniform resource identifier:

```
https://your_hmc_ip_address:12443/rest/api/pcm/preferences
```

To enable data collection for one or more servers, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .

   a) Click **Console Settings**. The **Console Settings** page is displayed.

   b) In the **Performance Settings** area, click **Change Performance Monitor Settings**.

   Note: Alternatively, if you try to launch the Performance and Capacity Monitor function for a system on which data collection is disabled, a message is displayed in the Current Resource Utilization area of the Settings for Performance Monitoring page. To enable data collection for the required server, set **Collection** to **On**.

2. Specify the number of days for which you want to store performance data by typing in a number in the range 1 - 366. Otherwise, click the up or down arrows next to **Number of days to store performance data** under **Performance Data Storage**.

   Note: By default, the HMC is set to store data for 180 days. The maximum number of days for which it can store data is 366.

3. Click the toggle switch in the **Collection** column, next to the name of the server, for which you want to collect data. Otherwise, click **All On** to enable data collection for all of the servers that the HMC manages.

   Note: If you request to monitor a number of managed servers that exceed the maximum number of managed servers that can be monitored by the HMC, the HMC displays an error.

4. Click **OK** to apply the changes and close the window.

   The **HMC Management** topic pane is displayed in the main window. You can now review the collected data by accessing the Performance and Capacity Monitor home page.

Accessing the Performance and Capacity Monitor home page

After you enable data collection for a server, the Performance and Capacity Monitor function plots the data in graphs and summarizes the information in tables. You can view the graphs and tables from the Performance and Capacity Monitor home page, which is accessible from the Hardware Management Console (HMC).

About this task

To access the Performance and Capacity Monitor home page, complete the following steps:

Procedure

In the navigation pane, click the **Resources** icon .

a) Click **All Systems**. The All Systems page is displayed.

b) Select the server for which you want to view the performance data.

c) Click **Actions**.

d) Select **View Performance Dashboard**.

The Performance and Capacity Monitor home page is displayed with the information for that system.

The Performance and Capacity Monitor home page

The Performance and Capacity Monitor home page contains charts and graphs that represent the data that is collected from the server.

The home page is divided into the following sections:

- The **Current Resource Utilization** graphs appear at the upper section of the Performance and Capacity Monitor home page. These graphs indicate current processor usage and memory assignment as a portion of available capacity. The Virtual I/O network traffic and storage bandwidth graphs indicate their current usage against their historical maximum bandwidth consumption that was recorded on that HMC after the performance monitoring function was enabled. You can change the chart auto-update interval. To view a larger version of a graph, click the icon that resembles a magnifying glass with a plus sign. To view help information for a graph, click the question mark icon.
- The **Views** topic pane appears on the right side of the home page and includes a list of server resources for which you can view performance data. The views include **Server Overview**, **Processor Utilization Trend**, **Memory Utilization Trend**, **Network Utilization Trend**, **Storage Utilization Trend**, and **SR-IOV Port Counters**.
- The details section occupies the remaining space on the Performance and Capacity Monitor home page. The details section displays the graphs and charts that are associated with the view that you selected from the **Views** topic pane.

Changing Performance and Capacity Monitor home page settings

You can change the time interval settings for the graphs on the Performance and Capacity Monitor home page.

Changing the automatic update frequency of Current Resource Utilization graphs

The **Current Resource Utilization** graphs default to an auto-update value of 1 minute; however, you can specify a longer time interval, if required.

To change the duration of time between updates, complete the following steps:

1. In the upper-right corner of the **Current Resource Utilization** section, click the menu next to **Auto-update in**.
2. Select one of the following preset values: **1 minute**, **5 minutes**, **10 minutes**, or **15 minutes**.

   The data in the graphs refreshes according to the time interval you chose.

Changing the time interval of the data that is displayed in the details pane

The home page defaults to a 4-hour time interval for the data in the details section. However, you can specify a longer time interval. You can also specify custom dates and times. The minimum time interval is 4 hours, and the maximum time interval is one year from the current date and time.

The details section refreshes and displays the updated content that is based on the time interval you choose. After the window refreshes, the latest entry for the data ends with the current time. The Performance and Capacity Monitor displays the data in this interval, every time you refresh the view, unless you change the interval again.

To change the time interval, complete the following steps:

1. Click the menu in the upper-right corner of the details section.
2. Select one of the following preset values: **Last 4 Hours**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**. Otherwise, select **Custom**.

If you selected **Custom**, a window is displayed. Continue to the next step.

3. Specify the date and time information in the **Start Date** and **End Date** fields, or click the calendar icon to choose the start date from the calendar.
4. Click **OK** to apply your changes.

Note: If you change the time interval of one view, then the interval change applies only to that view. For example, if you change the time interval for the **Server Overview** page to **Last Week**, the time interval for the Processor Trend view remains at **Last 4 Hours**.

Current Resource Utilization graphs

The upper section of the Performance and Capacity Monitor home page includes the **Current Resource Utilization** information. These graphs depict how the system's processor usage, memory assignment, virtual I/O network, and storage traffic compare against the available capacities or against their maximum historic highs. If the overall server usage is consistently high, activate more processors, move workloads to other servers, or buy more servers, processors, or memory.

The **Current Resource Utilization** information includes the following graphs: **Processor Usage/Peak**, **Memory Assignment**, **Network Traffic**, and **Storage Traffic**.

Processor Usage/Peak

The **Processor Usage/Peak** graph shows the average processor usage that is measured in processor cores and is represented by the blue horizontal bar. The black vertical bar indicates the maximum number of processors that the system used during the most recent monitoring period. The gray shading indicates the percentage of total active physical processors used. Light gray indicates that zero to 50% of the available processors were used. Medium gray indicates that 51% - 90% of the available processors were used. Dark gray indicates that 91% - 100% of the available processors were used. This graph indicates how the current and recent peak utilization compares against the total number of processors available on the server.

Click the **Click to Enlarge** button in the **Processor Usage/Peak** area. The table displays information about the Installed, Activated, Allocated, Available, Utilized, Peak, and Minimum usage data of processor usage.

Memory Assignment

The **Memory Assignment** graph shows the average memory assignment that is measured in MB or GB and is represented by the blue horizontal bar. The black vertical bar indicates the maximum amount of memory that the system used. The gray shading indicates the percentage of total active memory used. Light gray indicates that 0% - 50% of active memory was used. Medium gray indicates that 51% - 90% of active memory was used. Dark gray indicates that 91% - 100% of active memory was used. This graph indicates how the current and recent peak utilization compares against the total amount of memory available on the server.

Click the **Click to Enlarge** button in the **Memory Assignment** area. The table displays information about the Installed, Activated, Allocated, Available, Utilized, Peak, and Minimum memory assignment. You can also view information about the Assigned Virtual Persistent Memory. The persistent memory is available only when the Hardware Management Console (HMC) is at Version 9.1.940, or later, and when the firmware is at level FW940, or later. The persistent memory is a virtualization feature in which persistent memory volumes are created by using the existing DRAM (Dynamic Random Access Memory), where data persistence is maintained across applications, and the operating systems even when the logical partition has been restarted.

Network Traffic

The **Network Traffic** graph shows the average amount of traffic (measured in KB/s or GB/s and is represented by the blue horizontal bar) that flows through the network adapters that are assigned to the Virtual I/O Server. The gray shading indicates the maximum amount of traffic that was measured in the time that elapsed since the console was started. The black vertical bar indicates the maximum amount of network bandwidth that the system used. This graph indicates how the average network traffic compares against the maximum amount of network bandwidth that the system used.

Click the **Click to Enlarge** button in the **Network Traffic** area. The table displays information about traffic utilization that is measured per second (measured in KB/s or GB/s), the maximum amount of network bandwidth that the system used, the number of Virtual I/O Servers, and the number of Physical Adapters.

Storage Traffic

The **Storage Traffic** graph shows the average amount of traffic (measured in KB/s or GB/s and is represented by the blue horizontal bar) that is processed through the storage adapters and assigned to the Virtual I/O Server. The gray shading indicates the maximum amount of traffic that was measured in the time that elapsed since the console was started. The black vertical bar indicates the maximum amount of storage I/O bandwidth that the system used. This graph indicates how the average storage traffic compares against the maximum amount of storage I/O bandwidth that the system used.

Click the **Click to Enlarge** button in the **Storage Traffic** area. The table displays information about storage utilization that is measured per second (measured in KB/s or GB/s), maximum amount of storage I/O bandwidth that the system used, the number of Virtual I/O Servers, and the number of Physical Adapters.

The Server Overview section

The **Server Overview** section contains graphs that summarize data from virtualized server resources. This information indicates how physical processor and memory resources are allocated among the partitions on your server. Additionally, the information indicates whether partitions are using more or less than their entitled capacity for these resources. By default, the Performance and Capacity Monitor displays the data in the Details section of the home page.

The **Server Overview** section includes two general graphs: **Capacity Distribution by Processor** and **Capacity Distribution by Memory**. These graphs indicate general information about the capacity distribution for processors and memory.

The **Top Resource Consumers** graph displays information about partitions, Virtual I/O Servers, or processor pools. The **Resource Utilization** table shows detailed information about individual partitions, such as the number of processor cores and the amount of memory.

By default, the graphs and table show data that was collected for the previous 4 hours. For more information on displaying data for a longer amount of time, see "Changing Performance and Capacity Monitor home page settings" on page 3.

Note: When the firmware is at level 7.8, or later, and VIOS version is at 2.2.3, or later, you can view all the performance metrics. For more information about the limitations of the performance monitor metrics based on the firmware level and VIOS version, see HMC Integrated Performance Monitor Metrics based on Firmware and VIOS level.

The Capacity Distribution by Processor graph

The **Capacity Distribution by Processor** graph shows the percentage and number of partitions whose processor usage is high, medium, or low relative to the partitions' entitled processor capacity. The Performance and Capacity Monitor designates processor utilization as high if the percentage is 91% or greater, medium if the percentage is in the range 50% - 90%, and low if the percentage is 50% or lower.

No additional configurations are available for this graph. However, you can view a more detailed version. For more information, see "Accessing and reviewing the Detailed Spread graphs" on page 6.

The Capacity Distribution by Memory graph

The **Capacity Distribution by Memory** graph displays the percentage and number of partitions whose memory is high, medium, and low capacity relative to the partitions' entitled memory capacity. The Performance and Capacity Monitor designates memory usage as high, if the percentage is 91% or greater, as medium, if the percentage is in the range of 50% - 90%, and low if the percentage is 50% or lower.

No additional configurations are available for this graph. However, you can view a more detailed version. For more information, see "Accessing and reviewing the Detailed Spread graphs" on page 6 for more information.

Accessing and reviewing the Detailed Spread graphs

The **Detailed Spread** graphs provide details about the partition metrics that are shown in the **Capacity Distribution by Processor** and by Memory graphs. The graphs show dots that represent individual partitions whose current processor usage (vertical axis) is plotted against entitlement (horizontal axis). The diagonal lines have slopes of 0.5, 0.9, and 1.0, which represent usage, relative to entitlement of 50%, 90%, and 100%. If a partition is positioned above the 1.0 line, the partition is using more than 100% of its entitled capacity. You can hover your mouse pointer over a marker on the graph to view the name of the associated partition.

About this task

To access and review the Detailed Spread graphs, complete the following steps:

Procedure

1. On the Performance and Capacity Monitor home page, in the Server Overview section click **Show Detailed Spread**.

   The **All Partitions Spread** window is displayed.
2. Click **More Graphs** to switch between the **Processor Usage vs Entitlement** and the **Memory Usage vs Assigned** views.

The Top Resource Consumers graph

The **Top Resource Consumers** graph displays up to 10 partitions or Virtual I/O Servers that are using the highest number of units of the resource you chose.

Each vertical line represents a single partition, Virtual I/O Server, or processor pool. The top of each vertical line corresponds to the maximum number of resource units that are used, and the bottom of each line represents the minimum number of resource units.

The horizontal lines that bisect the vertical lines represent the average utilization of the resource.

The Resource ID appears along the bottom of the graph directly below the vertical line of the partition, Virtual I/O Server, or processor pool that the line represents. You can hover your mouse pointer over this area of the graph to view numeric values for minimum, maximum, and average utilization.

Changing the Top Resource Consumers graph

The **Top Resource Consumers** graph defaults to show up to 10 partitions that are using the most processors. However, you can change the graph to show the 10 partitions that are using the most memory, network, or storage resources. You can also choose to view the 10 highest processor pools.

About this task

To change the Top Resource Consumers graph to another graph, complete the following steps:

Procedure

1. On the Performance and Capacity Monitor home page, click **More Graphs**.
2. Select one of the following options:

   - **Partitions**
   - **Virtual I/O Servers**
   - **Processor Pools**

   If you chose **Partitions** or **Virtual I/O Servers**, continue with the next step. If you chose **Processor Pools**, there are no additional selections; the graph refreshes and shows the top 10 partitions that are using the processor pools.

3. Select one of the following resources:

   - **Processor**
   - **Memory**
   - **Network**
   - **Storage**

   The graph refreshes and shows the top 10 partitions or Virtual I/O Servers that are using the resource that you chose.

   Note: If you have fewer than 10 partitions or Virtual I/O Servers, the graph shows all of them.

The Resource Utilization table

The **Resource Utilization** table shows the amount of server resources, such as processor or memory, that is used by each partition. You can sort and filter the table. You can click the partition names in the Resource Utilization table to view the configuration information about the partition.

The **Resource Utilization** table also shows information about the Assigned Virtual Persistent Memory. The persistent memory is available only when the Hardware Management Console (HMC) is at Version 9.1.940, or later, and when the firmware is at level FW940, or later.

Sorting the Resource Utilization table

You can sort the **Resource Utilization** table by clicking the up or down chevron next to the name of the column that you want to sort. You might choose to sort the columns so that you can view the entries ranking from the lowest to the highest, or from the highest to the lowest.

You can select which columns are displayed in the **Resource Utilization** table. To change which columns are displayed, click the arrow in the header row of the table.

Filtering the Resource Utilization table

You can search for specific entries such as the partition name within the table. The search shows all table rows that contain text in any cell that matches the filter text.

The Processor Utilization view

The **Processor Utilization** view includes historical data and trends that reflect the usage of virtualized or shared processors over time. A graph shows the processor utilization on the physical server. Another aggregated graph shows the usage per resource, which includes the system firmware, Virtual I/O Servers, and client partitions. The table lists more detailed information about averages and trends.

You can access this view by clicking **Processor Utilization Trend** in the **Views** window.

The **Processor Utilization** view includes a trend graph. You can change the graph options to show processor utilization for the server and aggregated levels.

The **Resource Utilization** table shows detailed information for individual partitions and pools, such as the number of entitled and used units.

By default, the graphs and tables show data that was collected for the previous 4 hours. For more information about displaying data for a longer amount of time, see "Changing Performance and Capacity Monitor home page settings" on page 3.

Processor trend graphs

The Performance and Capacity Monitor home page includes trend graphs that show processor utilization data that is plotted against a default time interval of 4 hours.

By default, the trend graph shows server level data; however, you can change the view to display aggregated level data. To switch from one view to another view, click **More Graphs** and choose **Server Level Utilization** or **Aggregated Level Utilization**.

Processor trend graph: Server Level Utilization view

The **Server Level Utilization** view indicates the number of processors that the server is using at the times indicated along the horizontal axis. The lower shaded area represents the total number of activated physical processors on the server, and the upper shaded area indicates how many more processors are available for activation. The line shows how the total processor usage on the physical server varies over the selected time period, in comparison with the available processor capacity.

Processor trend graph: Aggregated Level Utilization view

The **Aggregated Level Utilization** view shows the total number of processors that the server is using. You can compare whether the processors are being used by the system firmware, Virtual I/O Servers, or client partitions by viewing the shading for each processor.

Processor breakdown tables

The processor breakdown tables list information that is based on partitions or pools over the selected time period. The following breakdown tables are available: **Breakdown by Partitions** and **Breakdown by Pools**.

The processor Breakdown by Partitions table

The **Breakdown by Partitions** table shows processor utilization data for logical partitions. Each row indicates whether a partition is using dedicated or shared processor resources. If the partition is using shared processor resources, the **Pool** column indicates the shared processor pool from which the resources are drawn.

In addition, you can view the number of processors or processor pools that the partition is entitled to use, is using, and the maximum number the partition used. The **Usage Trend** column shows the overall usage trend for the logical partition for the time interval you selected.

The table lists the total number of partitions for your system. The **Donated Units** column indicates whether the partition is donating unused processor resources to its shared processor pool. The **Dispatch Wait Time** column indicates the mount of time for which partitions are waiting for processor resources to be available.

The processor Breakdown by Pools table

The **Breakdown by Pools** table shows processor utilization within individual processor pools. You can view the total processor entitlement of all partitions that use resources from the pool and the number of processors the pool borrowed. You can also view the number of processors the pool is using and the maximum number of processors the pool used.

The **Usage Trend** column shows a high-level trend view for an individual pool.

Sorting the processor breakdown tables

You can sort the table by clicking the up or down chevron next to the name of the column that you want to sort. You might choose to sort the columns so that you can view the entries, ranking from the lowest to the highest, or vice versa.

You can select which columns are displayed in the **Processor Breakdown** tables. To change which columns are displayed, click the arrow in the header row of the table.

Filtering the processor breakdown tables

You can search for specific entries such as the partition name within the table. The search shows all table rows that contain text in any cell that matches the filter text.

The Memory Utilization view

The **Memory Utilization** view includes historical data and trends that reflect the amount of dedicated memory that is allocated, or shared among logical partitions, over time. The graph shows memory usage that is divided by the total, allocated, and assigned usage. The table lists more detailed information about averages and trends.

You can access this view by clicking **Memory Utilization Trend** in the **Views** window.

The **Memory Utilization** view includes a trend graph. You can change the graph options to show server level, aggregated level, or Active Memory Sharing (AMS) level memory utilization.

The **Resource Utilization** table shows detailed information for individual partitions. The information includes the amount of memory that the firmware is using and the amount of memory in the shared memory pool that partitions are entitled to use.

By default, the graphs and table show data that was collected over the last 4 hours. To display data for a longer amount of time, refer to "Changing Performance and Capacity Monitor home page settings" on page 3.

Memory trend graphs

The Performance and Capacity Monitor home page includes trend graphs that show memory utilization data that is plotted against a default time interval of 4 hours.

By default, the trend graph shows server level data; however, you can change the view to display aggregated or Active Memory Sharing (AMS) level data. To switch from one view to another view, click **More Graphs** and choose **Server Level Utilization**, **Aggregated Level Utilization**, or **AMS Level Utilization**.

Memory trend graph: Server Level Utilization view

The **Server Level Utilization** view shows the memory usage for the server. Shaded areas indicate the amount of memory that is assigned to the server, the amount of memory that is allocated for use by the server, and the total memory available for use. You can compare the shaded areas to determine whether you maximized the memory allocation for your server.

Memory trend graph: Aggregated Level Utilization view

The **Aggregated Level Utilization** view shows the collective memory usage for the partitions on that server. Shaded areas indicate the amount of memory that is allocated to the system firmware, the amount of memory that is used by all Virtual I/O Servers, and the amount of memory available for all partitions. You can compare the trend lines to determine whether you allocated more memory or less memory for the partitions on your server.

Memory trend graph: AMS Level Utilization view

The **AMS Level Utilization** view shows the amount of memory that is used by Active Memory Sharing (AMS). The shaded area indicates the amount of memory that is used by the system firmware at the times that are displayed along the horizontal axis. You can review this information periodically to determine whether your system benefits from using memory from Active Memory Sharing. If you are not using Active Memory Sharing, this information is not available.

Memory breakdown table

The memory **Breakdown by Partitions** table lists information that is based on partitions over the selected time period. The memory breakdown table is displayed at the bottom of the main window.

The Memory Breakdown by Partitions table

The **Breakdown by Partitions** table shows memory utilization data for logical partitions. Each row indicates whether the partition is configured for dedicated or shared access to memory resources. In addition, you can view the size of the memory available, the amount of memory that is assigned, and the maximum amount of memory that is assigned to the memory pool for that partition. The **Assigned Trend** column shows the overall usage trend for the assigned memory over the time interval you selected. The **Breakdown by Partitions** table also lists the total number of partitions for that system.

The **Breakdown by Partitions** table also shows information about the Assigned Virtual Persistent Memory. The persistent memory is available only when the Hardware Management Console (HMC) is at Version 9.1.940, or later, and when the firmware is at level FW940, or later.

Sorting the memory breakdown table

You can sort the table by clicking the up or down chevron next to the name of the column that you want to sort. You might choose to sort the columns so that you can view the entries, ranking from the lowest to the highest, or vice versa.

You can select which columns are displayed in the **Memory Breakdown** table. To change which columns are displayed, click the arrow in the header row of the table.

Filtering the memory breakdown table

You can search for specific entries such as the partition name within the table. The search shows all table rows that contain text in any cell that matches the filter text.

The Network Utilization view

The **Network Utilization** view includes historical data and trends that show how logical partitions use physical network resources, or virtual local area network resources, over time. This view contains a graph that shows the network traffic per Virtual I/O Server. The table lists more detailed information about the averages and trends.

You can access this view by clicking **Network Utilization Trend** in the **Views** window.

The **Network Utilization** view includes a trend graph.

The **Resource Utilization** table shows detailed information for individual partitions and network bridges, such as the amount of traffic that is processed by the physical resources.

By default, the graphs and tables show data that was collected for the previous 4 hours. To display data for a longer amount of time, refer to "Changing Performance and Capacity Monitor home page settings" on page 3.

Network trend graphs

The Performance and Capacity Monitor home page includes a trend graph that shows network utilization data that is plotted against a default time interval of 4 hours.

By default, the trend graph shows network bridge level data; however, you can change the view to display Single root I/O virtualization (SR-IOV) adapters traffic data. To switch from one view to another view, click **More Graphs** and choose **Network Bridges Traffic** or **SR-IOV Adapters Traffic**.

Network trend graph: Network Bridges Traffic trend view

The **Network Bridges Traffic** view shows the traffic that is flowing over network bridges at the times indicated along the horizontal axis. The shaded areas indicate the amount of internal virtual traffic (measured in GB per second) that is tagged by a Virtual I/O Server and is processed by shared Ethernet adapters. The dotted line indicates the amount of physical traffic that is routed to a physical NIC for sharing outside of the virtual network. You can compare the shaded areas to determine how much virtual traffic is sent to one Virtual I/O Server versus another. Similarly, you can view the dotted line to compare the amount of physical traffic versus the amount of virtual traffic.

Network breakdown tables

The network breakdown tables list information about network traffic for the selected time period. The following breakdown tables are available: **Breakdown by Partitions** and **Breakdown by Network Bridges**.

The network Breakdown by Partitions table

The **Breakdown by Partitions** table shows network traffic data for logical partitions. Each row indicates the ID of the network bridge with which that partition is associated. The rows also indicate the number of Virtual I/O Servers that are associated with the partition, and the amount of virtual and physical traffic that is processed through the partition. The **Traffic Trend** column shows the overall network traffic for the logical partition for the time interval you selected.

Click a network bridge ID to display network traffic information, such as the number of packets that were sent and received and the speeds at which packets were sent or received, for the bridge.

The network Breakdown by Network Bridges table

The **Breakdown by Network Bridges** table shows the network traffic for network bridges. Each row indicates the name of the network bridge, the number of partitions that are sending traffic across that bridge, the name of the Virtual I/O Server that hosts the network bridge, and the amount of virtual and physical traffic that is processed through the bridge. The **Traffic Trend** column shows the overall network traffic on the network bridge for the time interval you selected.

Click a network bridge ID to display network traffic information, such as the number of packets that were sent and received and the speeds at which packets were sent or received, for the bridge.

Click one of the numbers in the **Partitions Using** column to view the names of the partitions that are using that network bridge.

Sorting the network breakdown tables

You can sort the table by clicking the up or down chevron next to the name of the column that you want to sort. You might choose to sort the columns so that you can view the entries, ranking from the lowest to the highest, or vice versa.

You can select which columns are displayed in the **network breakdown** table. To change which columns are displayed, click the arrow in the header row of the table.

Filtering the network breakdown tables

You can search for specific entries such as the partition name within the table. The search shows all table rows that contain text in any cell that matches the filter text.

The Storage Utilization view

The **Storage Utilization** view includes historical data and trends that show the amount of physical storage I/O bandwidth that each Virtual I/O Server uses and allows logical partitions to use through Small Computer System Interface (SCSI) connections, for a specific duration of time. The data also shows the virtualized storage I/O bandwidth that the logical partitions use from logical ports that are provided by an N_Port ID Virtualization (NPIV) adapter. The table lists detailed information about averages and trends.

You can access this view by clicking **Storage Utilization Trend** in the **Views** window.

The **Storage Utilization** view includes a trend graph. You can change the graph options to show **vSCSI Adapter usage** or **NPIV traffic**.

The **Resource Utilization** table shows detailed information for individual partitions and physical adapters, such as the total traffic used.

By default, the graphs and tables show data that was collected for the previous 4 hours. For more information about displaying data for a longer amount of time, see "Changing Performance and Capacity Monitor home page settings" on page 3.

Storage trend graphs

The Performance and Capacity Monitor home page includes trend graphs that show storage utilization data that is plotted against a default time interval of 4 hours.

By default, the trend graph shows data for virtual Small Computer System Interface (SCSI) adapters; however, you can change the view to display N_Port ID Virtualization (NPIV) traffic data. To switch from one view to another, click **Graph Options** and choose **vSCSI Adapaters Usage** or **NPIV Traffic**.

Storage trend graph: vSCSI Adapters Usage view

The **vSCSI Adapters Usage** view shows the I/O bandwidth for a Virtual I/O Server (VIOS) that is using physical storage space on SCSI adapters at the times indicated along the horizontal axis. Each of the shaded areas represents one VIOS. You can compare the shaded areas with one another to determine which VIOS is using the most storage bandwidth, and you can compare individual VIOS usage against the total usage.

Storage trend graph: NPIV Traffic view

The **NPIV Traffic** view shows the I/O bandwidth for a VIOS that is using physical storage space through logical ports that are provided by an NPIV adapter at the times indicated along the horizontal axis. Each of the shaded areas represents one VIOS. You can compare the shaded areas with one another to determine which VIOS is using the most storage bandwidth, and you can compare individual VIOS usage against the total usage.

Storage breakdown table

The storage breakdown tables list information that is based on partitions or physical Fibre Channel (FC) adapters over the selected time period. The following breakdown tables are available: **Breakdown by Partitions** and **Breakdown by Pools**.

The storage Breakdown by Partitions table

The **Breakdown by Partitions** table shows the amount of traffic that is passing through the physical storage adapter that is associated with the logical partition. Each row indicates the name of the Virtual I/O Server and Virtual Host that is associated with the partition. The **Traffic Trend** column shows the overall traffic trend for the physical adapter over the time interval you selected.

The table lists the total number of partitions for that system.

The storage Breakdown by Physical FC table

The **Breakdown by Physical FC** table shows the amount of traffic that is passing through the physical storage adapter that is associated with the Physical FC. Each row indicates the name of the Virtual I/O Server and Virtual Host that is associated with the Physical FC. The **Traffic Trend** column shows the overall traffic trend for the physical adapter over the time interval you selected.

Sorting the storage breakdown tables

You can sort the table by clicking the up or down chevron next to the name of the column that you want to sort. You might choose to sort the columns so that you can view the entries, ranking from the lowest to the highest, or vice versa.

You can select which columns are displayed in the **storage breakdown** tables. To change which columns are displayed, click the arrow in the header row of the table.

Filtering the storage breakdown tables

You can search for specific entries such as the partition name within the table. The search shows all table rows that contain text in any cell that matches the filter text.

Viewing SR-IOV port counters

You can view SR-IOV port counters in the Hardware Management Console (HMC) version 8.7.0, or later. The **SR-IOV Port Counters** page displays the details of logical ports and physical ports that are configured for a selected SR-IOV adapter. You can use the **SR-IOV Port Counters** page to view port counters for a logical port or a physical port that is configured for a selected SR-IOV adapter.

About this task

To view the SR-IOV port counters, by using the Hardware Management Console (HMC), complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**.

   The **All Systems** page is displayed.
3. Select the server for which you want to view the performance data.
4. Click **Actions**.
5. Select **View Performance Dashboard**.
6. In the **Views** topic pane, click **SR-IOV Port Counters**.

   The **SR-IOV Port Counters** page is displayed. The **SR-IOV Adapters** list displays the details of the SR-IOV adapters that are configured for the selected system.
7. From the **SR-IOV Adapters** list, select an SR-IOV adapter.

   The status, mode, owner, maximum logical ports, and the configured logical ports of the selected SR-IOV adapter are displayed.
8. Select **Physical Ports** or **Logical Ports** to view the list of physical or logical ports that are configured for the SR-IOV adapter.

   - If you want to view the list of physical ports that are configured for the SR-IOV adapter, select **Physical Ports**. The **Physical Ports** table is displayed, with the details about the physical ports such as ID, location code, type, link status, label, and sub label of the port.

     Note: If the selected SR-IOV adapter does not have any physical ports that are attached to it, the **Physical Ports** table does not display any details about the physical ports.

     To view the list of port counters, complete the following steps:

     a. From the **Physical Ports** table, select a physical port to view the list of port counters. The **Port Counters** table is displayed, with the name and value of port counters for the selected physical port.

     b. Click **Reset Statistics** to reset the port counter statistics of the selected physical port.

- If you want to view the list of logical ports that are configured for the SR-IOV adapter, select **Logical Ports**. The **Logical Ports** table is displayed, with the details about the logical ports such as adapter ID, physical port ID, location code, type, partition, and connected partition for the port.

  Note: If the selected SR-IOV adapter does not have any logical ports that are attached to it, the **Logical Ports** table does not display any details about the logical ports.

  To view the list of port counters, complete the following steps:

  a. From the **Logical Ports** table, select a logical port to view the list of port counters. The **Port Counters** table is displayed, with the name and value of port counters for the selected logical port.
  b. Click **Reset Statistics** to reset the port counter statistics of the selected logical port.

Troubleshooting the Performance and Capacity Monitor

Review the common troubleshooting issues and their solutions.

How can I determine whether performance data is being collected?

The Performance and Capacity Monitor home page includes a **Data Collection** status indicator on the home page. If the status is **On**, the Performance and Capacity Monitor function is collecting data from that server. If the status is **Off**, the Performance and Capacity Monitor function is not collecting data from that server. For more information about collecting data from your system, see "Enabling data collection" on page 2.

What permission do I need to view the managed system utilization data?

You must have **List Utilization Data** access permission for the managed system to view the performance data for that server. For more information about user roles and permissions, see HMC tasks, user roles, IDs, and associated commands .

What happens if I power off the system while Performance and Capacity Monitor is still enabled?

If you power off a PCM-enabled system, Performance and Capacity Monitor does not get disabled automatically. Instead, the Performance and Capacity Monitor data collection that happens in the background is stopped. However, the status of Performance and Capacity Monitor for the PCM-enabled system is still displayed as **On** in the graphical user interface (GUI) and in Representational State Transfer (REST). When the system is powered on again, Performance and Capacity Monitor restarts data collection, as usual.

Why was data not collected for my server even though I enabled data collection?

You can enable data collection for servers that are in any state. However, the Performance and Capacity Monitor collects data in the Hardware Management Console (HMC) only when the server is in running or in the operational state. The Performance and Capacity Monitor automatically disables collection, if the server is not in running or operational state for 30 minutes or longer.

Why does the home page not display data even though I enabled data collection?

If you access the Performance and Capacity Monitor home page before the initial data is collected, the Performance and Capacity Monitor displays a status message. The status message indicates that data is not yet available and recommends that you go to the home page again later. The initial time that is required to collect the information is about 15 minutes.

Why is my physical adapter not displayed under the PCM data?

If a storage adapter is not connected to a device, that adapter does not appear under the Performance and Capacity Monitor data, as there is no utilization for that adapter.

Why are the Performance and Capacity Monitor graphs not displayed, and instead I see only the "Fetching PCM Data" message?

You must clear the browser cache and cookies, and then try again.

Why is the home page not showing data for the entire length of time I chose?

The Performance and Capacity Monitor home page can show only the amount of data that the server stored since you enabled data collection. For example, if you want to collect data for 250 days and if you immediately access the home page, you can view only the data that represents the minute or minutes that passed since you enabled data collection.

In addition, the maximum number of days for which the Performance and Capacity Monitor collects data is 366. As a result, the Performance and Capacity Monitor shows a maximum of 366 days of data only.

Why do I see gaps in the data that is displayed in the collection graphs?

If you disable the data collection and re-enable it, or if the server stopped collecting data because the server stopped running, or it is no longer operational, the Performance and Capacity Monitor shows gaps that represent the missing time intervals.

Can I see utilization data after I disable the data collection?

Yes, the Performance and Capacity Monitor maintains utilization data after data collection is disabled. You can view the historic data from the Performance and Capacity Monitor home page of your server. For more information, see "Accessing the Performance and Capacity Monitor home page" on page 2.

Why do I receive a message that indicates that network or storage resources are not available to display?

If you dedicate network and storage resources to a single partition on your server, network and storage utilization data is not available. Network and storage utilization data shows how each of the partitions on your server is using network and storage resources that are managed by Virtual I/O Servers. You can compare the data among partitions to determine whether a partition is overloaded or under-used. However, if a single partition is entitled to dedicated network and storage resources, there is no data to compare. In addition, you can also check whether you have the required Virtual I/O Server version. The Performance and Capacity Monitor requires Virtual I/O Server Version 2.2.3, or later to display the Network and Storage data.

Why do I see only a single partition or Virtual I/O Server listed in the Top Resource Consumers graph?

The **Top Resource Consumers** graph displays up to 10 partitions or Virtual I/O Servers that are using the highest number of units of the resource you chose. However, if you dedicate all of your resources to a single partition or Virtual I/O Server, no other partitions or servers can compete for the resources. As a result, only the partition or Virtual I/O Server for which you dedicated all resources is displayed in the **Top Resource Consumers** graph.

Similarly, if you have fewer than 10 partitions or Virtual I/O Servers, the **Top Resource Consumers** graph includes a vertical line for each of your partitions or Virtual I/O Servers. A maximum of 10 partitions or Virtual I/O Servers are included in the **Top Resource Consumers** graph. If fewer than 10 partitions or Virtual I/O Servers exist, all the partitions or Virtual I/O Servers are displayed.

Disabling data collection

The Performance and Capacity Monitor function captures data only for the servers for which you enabled data collection. However, you can disable data collection if you no longer need performance and capacity monitoring information for that server.

About this task

To disable data collection, complete the following steps:

Procedure

1. In the navigation pane, click the **HMC Management** icon .
2. Click **Console Settings**. The **Console Settings** page is displayed.
3. In the **Performance Settings** area, click **Change Performance Monitor Settings**.

   Note: Alternatively, if you try to launch the Performance and Capacity Monitor function for a system on which data collection is disabled, a message is displayed in the Current Resource Utilization area of the **Settings for Performance Monitoring** page. To enable data collection for the required server, set **Collection** to **On**.
4. Click the toggle switch in the **Collection** column next to the name of the server for which you want to disable data collection or click **All Off** to disable data collection for all of the servers in your environment.
5. Click **OK** to apply the changes and close the window.

   The **HMC Management** content is displayed in the main window.

Exporting data

The Export Data option exports the Performance and Capacity Monitor (PCM) data metrics that are collected for the specified time. You can export the Performance and Capacity Monitor data metrics that is displayed in the dashboard into a folder on your local system.

About this task

You can export data metrics by accessing the HMC.

To export data metrics for one or more servers, complete the following steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
   a) Click **All Systems**. The All Systems page is displayed.
   b) Select the server for which you want to view the performance data.
   c) Click **Actions**.
   d) Select **Performance Data Collection** > **Export Data**.

   The Performance and Capacity Monitor home page is displayed with the information for that system.
2. In the upper-right corner of the Performance and Capacity Monitor section, click the **Data Collection** menu.
3. Click **Export Data**.

   The **Export Data** page is displayed.

4. The data collection can be turned on or turned off using the toggle switch in the **Data Collection** menu.
5. Select the feed by which you want to export the performance metrics.

   PCM metrics have the following rollup or aggregation frequency and retention period:

   - Tier 0 level - The aggregation frequency is 30 seconds and the retention period is 2 hours.
   - Tier 1 level - The aggregation frequency is 5 minutes and the retention period is 24 hours.
   - Tier 2 level - The aggregation frequency is 2 hours and the retention period is 7 days.
   - Tier 3 level - The aggregation frequency is 24 hours and the retention period is 180 days.

   When you select the feed as **By Source**, multiple data files that contain overall resource level data are exported for each managed system, logical partition, and Virtual I/O Server (VIOS). When you select the feed type as **By Tier**, the maximum tier level is calculated based on the time duration that is specified in start time stamp and end time stamp, and the data for corresponding level is exported.

   Note:

   If you select the feed type as **By Tier** and the export format is CSV, two files are created, one for the managed system and one for the logical partition, as compared to a single file in the JavaScript Object Notation (JSON) format.
6. Select the export format as either JavaScript Object Notation (**JSON**) or comma-separated values (**CSV**).

   The CSV file represents data of JSON file content in the CSV format.
7. Click the calendar icon to choose the **Start Date** and **End Date**.

   By default, the time stamp in the **Start Date** is set to 4 hours before the current time, and the current time is set as the time stamp in the **End Date**. You can choose to export data for this duration. Otherwise, to specify specific time interval, you must enter the time within the Performance and Capacity Monitor data retention period, which is 180 days by default. The time stamp of the last data export is displayed in gray.
8. Click **OK**. The **Confirm Download** dialog box is displayed with the name of the file that contains the exported data.
9. Click **OK** to download the exported data in a compressed format.
10. Depending on your browser setting, you can choose the destination folder in which the exported data must be saved.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Programming interface information

This Monitoring the Virtualization Environment publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.2, IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.4, and IBM Virtual I/O Server Version 3.1.1.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Beginning troubleshooting and problem analysis

Note

Before using this information and the product it supports, read the information in “Safety notices” on page v, “Notices” on page 131, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Power Systems servers that contain the POWER9™ processor and to all associated models.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
  - For AC power, disconnect all power cords from their AC power source.
  - For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected.
  - For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.

  - For racks with a DC power distribution panel (PDP), connect the customer’s DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.

- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements.
- Do not continue with the inspection if any unsafe conditions are present.
- Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.

⚠ **DANGER:**

- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect:

  1. Turn off everything (unless instructed otherwise).
  2. For AC power, remove the power cords from the outlets.
  3. For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source.
  4. Remove the signal cables from the connectors.
  5. Remove all cables from the devices.

  To Connect:

  1. Turn off everything (unless instructed otherwise).
  2. Attach all cables to the devices.
  3. Attach the signal cables to the connectors.
  4. For AC power, attach the power cords to the outlets.
  5. For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP.
  6. Turn on the devices.

  Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

⚠ **DANGER:** Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed..
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.

- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).

- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer's DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

CAUTION:

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.
  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

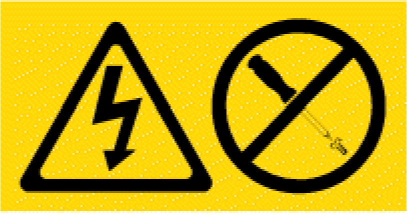

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

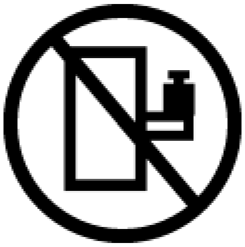

DANGER: Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.
- Before extending the rack to the installation position, read the installation instructions.
- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

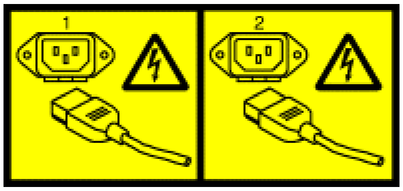

or

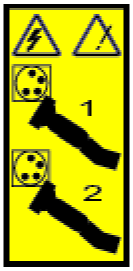

or

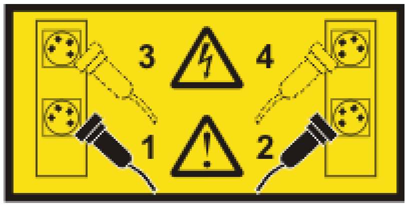

or

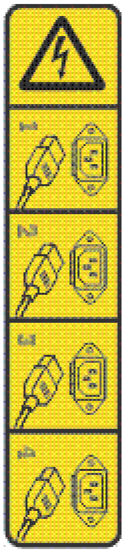

or

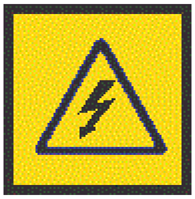

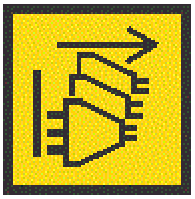

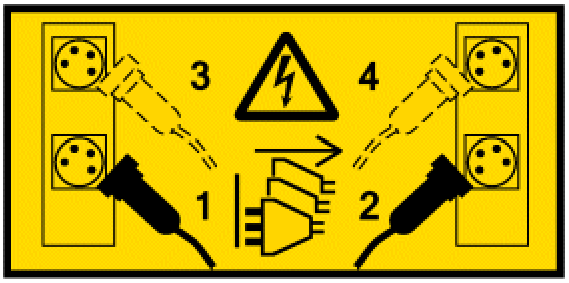

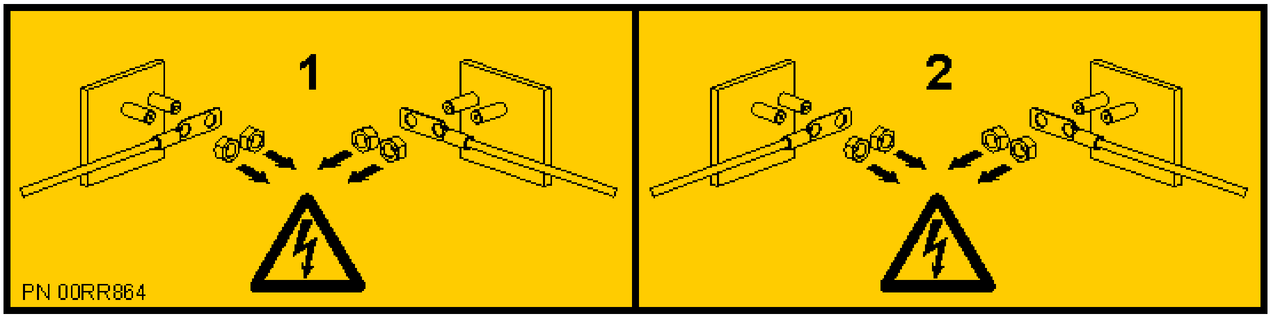

DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

CAUTION: A hot surface nearby. (L007)

(L008)

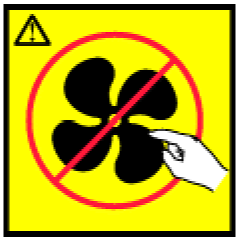

⚠ **CAUTION:** Hazardous moving parts nearby. (L008)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

⚠ **CAUTION:** This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

⚠ **CAUTION:** Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

⚠ **CAUTION:** This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

⚠ **CAUTION:** Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

⚠ **CAUTION:** The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)

⚠ **CAUTION:** Regarding IBM provided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.

- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.
- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four (4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.
- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely.

Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.

- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intrabuilding ports of this equipment are suitable for connection to intrabuilding or unexposed wiring or cabling only. The intrabuilding ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The ac-powered system does not require the use of an external surge protection device (SPD).

The dc-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The dc-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Beginning troubleshooting and problem analysis

This information provides a starting point for analyzing problems.

This information is the starting point for diagnosing and repairing servers. From this point, you are guided to the appropriate information to help you diagnose server problems, determine the appropriate repair action, and then perform the necessary steps to repair the server. A system attention light, an enclosure fault light, or a system information light indicates there is a serviceable event (an SRC in the control panel or in one of the serviceable event views) on the system. This information guides you through finding the serviceable event.

Beginning problem analysis

You can use problem analysis to gather information that helps you determine the nature of a problem encountered on your system. This information is used to determine if you can resolve the problem yourself or to gather sufficient information to communicate with a service provider and quickly determine the service action that needs to be taken.

If you are using this information because of a problem with your Hardware Management Console (HMC), see Managing the HMC.

To begin analyzing the problem, complete the following steps:

1. Do you have a direct indication of a hardware error (such as an automated email that notified you of a hardware error or a fault indicator on a system unit or expansion unit)?

   - **Yes:** Continue with the next step.
   - **No:** Go to "Detecting problems" on page 63.

2. How do you manage the system that is failing? If you do not know how the failing system is managed, ask the system administrator.

System management	Problem analysis
Hardware Management Console (HMC)	Go to the section "Hardware Management Console (HMC) problem analysis" on page 1.
Operating system (AIX®, Linux®, or IBM i)	Go to the problem analysis topic for your operating system. • If you are having a problem with an AIX or Linux system unit, go to "AIX and Linux problem analysis" on page 3. • If you are having a problem with an IBM i system unit, go to " IBM i problem analysis" on page 7.

Hardware Management Console (HMC) problem analysis

To perform beginning problem analysis on a system that is managed by Hardware Management Console (HMC), complete the following steps:

1. Is the management console functional and connected to the hardware?

   - **Yes:** Continue with the next step.
   - **No:** Start the management console and attach it to the system unit. Then return here and continue with the next step.

2. On the management console that is used to manage the system unit, complete the following steps:

   Note: If you are unable to locate the reported problem, and there is more than one open problem near the time of the reported failure, use the earliest problem in the log.

   a. In the navigation area, click the **Serviceability** icon, and then click **Serviceable Events Manager**. The Manage Serviceable Events window is displayed.
   b. In the Event Criteria area, for **Serviceable Event Status**, select **Open**. For all other criteria, select **ALL**, then click **OK**.

   Scroll through the log and verify that there is a problem with the status of Open to correspond with the failure.

   Do you find a serviceable event, or an open problem near the time of the failure?

   - **Yes:** Continue with the next step.
   - **No:** Contact your hardware service provider. **This ends the procedure.**

---

3. The reference code description might provide information or an action that you can take to correct the failure.

   Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action at this time.

   For more information about reference codes, see Reference codes.

   Was there a reference code description that enabled you to resolve the problem?

   - **Yes: This ends the procedure.**
   - **No:** Continue with the next step.

---

4. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

- If a FRU location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
- If an isolation procedure is listed for the reference code in the reference code lookup information, include the isolation procedure as a corrective action even if it is not listed in the serviceable event view or control panel.
- If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

From the Repair Serviceable Event window, complete the following steps:

a. Record the problem management record (PMR) number for the problem if one is listed.
b. Select the serviceable event from the list.
c. Click **Selected and View Details**.
d. In the Serviceable Event Details page, locate details such as the reference code and FRU list and record this information.
e. If a Problem Management Hardware (PMH) number was found for the problem on the Serviceable Event Overview panel, the problem has already been reported. If there was no PMH number for the problem, contact your service provider.

This ends the procedure.

AIX and Linux problem analysis

You can use this procedure to find information about a problem with your server hardware when service is managed by the AIX or Linux operating system.

Remember the following points while troubleshooting problems:

- Has an external power outage or momentary power loss occurred?
- Has the hardware configuration changed?
- Has system software been added?
- Have any new programs or program updates (including PTFs) been installed recently?

Before you use this procedure, ensure that you completed the steps in "Beginning problem analysis" on page 1.

After you review these considerations, complete the following steps:

1. Is the operating system operational?
   - **Yes:** Continue with the next step.
   - **No:** Go to step "11" on page 5.

2. Are any messages (for example, a device is not available or reporting errors) related to this problem displayed on the system console or sent to you in email that provides a reference code?

   Note: A reference code can be an 8 character system reference code (SRC) or a service request number (SRN) of 5, 6, or 7 characters, with or without a hyphen.

   - **Yes:** Continue with the next step.
   - **No:** Go to step "4" on page 4.

3. The reference code description might provide information or an action that you can take to correct the failure.

Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

For more information about reference codes, see Reference codes.

If the reference code description provides information to resolve the problem without replacing FRUs in the failing item list, complete the steps.

Were you able to resolve the problem?

- **Yes: This ends the procedure.**
- **No:** Continue with the next step.

4. Are you running the Linux operating system?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "6" on page 4.

5. To locate the error information in a system or logical partition that is running the Linux operating system, complete the following steps:

   Note: Before you proceed with this step, ensure that the diagnostics package is installed on the system.

   a. Log in as root user.
   b. At the command line, type `grep RTAS /var/log/platform` and press **Enter**.
   c. Look for the most recent entry that contains a reference code.

   Continue with step "8" on page 4.

6. To locate the error information in a system or logical partition that is running the AIX operating system, complete the following steps:

   a. Log in to the AIX operating system as root user, or use CE login. If you need help, contact the system administrator.
   b. Type `diag` to load the diagnostic controller, and display the online diagnostic menus.
   c. From the Function selection menu, select **Task selection**.
   d. From the Task selection list menu, select **Display previous diagnostic results**.
   e. From the Previous diagnostic results menu, select **Display diagnostic log summary**.

   Continue with the next step.

7. A display diagnostic log is shown with a time ordered table of events from the error log.

   Look in the T column for the most recent entry that has an S entry. Press **Enter** to select the row in the table and then select **Commit**.

   The details of this entry from the table are shown. Look for the SRN entry near the end of the entry and record the information that is shown.

   Continue with the next step.

8. Do you find a serviceable event or an open problem near the time of the failure?

   - **Yes:** Continue with the next step.
   - **No:** Contact your hardware service provider. **This ends the procedure.**

9. The reference code description might provide information or an action that you can take to correct the failure.

   Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

   For more information about reference codes, see Reference codes.

   Was there a reference code description that helped you to resolve the problem?

   - **Yes: This ends the procedure.**
   - **No:** Continue with the next step.

10. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

    - If a field-replaceable unit (FRU) location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
    - If an isolation procedure is listed for the reference code in the reference code lookup information, include it as a corrective action even if it is not listed in the serviceable event view or control panel.
    - If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

    From the Error Event Log view, complete the following steps:

    a. Record the reference code.
    b. Record the error details.
    c. Contact your service provider.

    This ends the procedure.

11. Details about errors that occur when the operating system is not running or when the operating system is now not accessible can be found in the control panel or in the Advanced System Management Interface (ASMI).

    Do you choose to look for error details by using ASMI?

    - **Yes:** Go to step “13” on page 5.
    - **No:** Continue with the next step.

12. At the control panel, complete the following steps.

    a. Press the increment or decrement button until the number 11 is displayed in the upper-left corner of the display.
    b. Press **Enter** to display the contents of function 11.
    c. Look for a reference code in the upper-right corner.

    Is a reference code displayed on the control panel in function 11?

    - **Yes:** Go to step “14” on page 6.
    - **No:** Contact your hardware service provider. **This ends the procedure.**

13. On the console that is connected to the ASMI, complete the following steps.

    Note: If you are unable to locate the reported problem, and there is more than one open problem near the time of the reported failure, use the earliest problem in the log.

a. Log in with a user ID that has an authority level as general, administrator, or authorized service provider.
b. In the navigation area, expand **System Service Aids** and click **Error/Event Logs**. If log entries exist, a list of error and event log entries is displayed in a summary view.
c. Scroll through the log under **Serviceable Customer Attention Events** and verify that there is a problem to correspond with the failure.

For information about the ASMI, see Managing the Advanced System Management Interface.

Do you find a serviceable event, or an open problem near the time of the failure?

- **Yes:** Continue with the next step.
- **No:** Contact your hardware service provider. **This ends the procedure.**

14. The reference code description might provide information or an action that you can take to correct the failure.

    Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

    For more information about reference codes, see Reference codes.

    Was there a reference code description that helped you to resolve the problem?

    - **Yes: This ends the procedure.**
    - **No:** Continue with the next step.

15. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

    - If a field-replaceable unit (FRU) location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
    - If an isolation procedure is listed for the reference code in the reference code lookup information, include the isolation procedure as a corrective action even if it is not listed in the serviceable event view or control panel.
    - If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

    To find error details on the control panel, complete the following steps:

    a. Press **Enter** to display the contents of function 14. If data is available in function 14, the reference code has a FRU list.
    b. Record the information in functions 11 through 20 on the control panel.
    c. Contact your service provider and report the reference code and other information.

    To find error details on the ASMI, complete the following steps from the Error Event Log view:

    a. Record the reference code.
    b. Select the corresponding check box on the log and click Show details.
    c. Record the error details.
    d. Contact your service provider.

    This ends the procedure.

IBM i problem analysis

You can use this procedure to find information about a problem with your server hardware when service is managed by the IBM i operating system.

If you experience a problem with your system or logical partition, try to gather more information about the problem to either solve it, or to help your next level of support or your hardware service provider to solve it more quickly and accurately.

This procedure refers to the IBM i control language (CL) commands that provide a flexible means of entering commands on the IBM i logical partition or system. You can use CL commands to control most of the IBM i functions by entering them from either the character-based interface or the IBM Navigator for i Web console. While the CL commands might be unfamiliar at first, the commands follow a consistent syntax, and IBM i includes many features to help you use them easily. The Programming navigation category in IBM i Knowledge Center includes a complete CL reference and a CL Finder to look up specific CL commands.

Remember the following points while troubleshooting problems:

- Has an external power outage or momentary power loss occurred?
- Has the hardware configuration changed?
- Has system software been added?
- Have any new programs or program updates (including PTFs) been installed recently?

To make sure that your IBM software was correctly installed, use the Check Product Option (CHKPRDOPT) command.

- Have any system values changed?
- Has any system tuning been done?

Before you use this procedure, ensure that you completed the steps in “Beginning problem analysis” on page 1.

After you review these considerations, follow these steps:

1. Is the IBM i operating system up and running?

   - **Yes:** Continue with the next step.
   - **No:** Go to step “19” on page 10.

2. Are you experiencing problems with the Operations Console?

   - **Yes:** See Troubleshooting Operations Console.
   - **No:** Continue with the next step.

3. Does the console show a Main Storage Dump Manager display?

   - **Yes:** Go to Copying a dump.
   - **No:** Continue with the next step.

4. Is the console that was in use when the problem occurred (or any console) operational?

   Note: The console is operational if a sign-on display or a command line is present. If another console is operational, use it to resolve the problem.

- **Yes:** Continue with the next step.
- **No:** Choose from the following options:
  - If your console does not show a sign-on display or a menu with a command line, go to Recovering when the console does not show a sign-on display or a menu with a command line.
  - For all other workstations, see the Troubleshooting navigation category in IBM i Knowledge Center.

---

5. Is a message related to this problem shown on the console?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "10" on page 8.

---

6. Is this a system operator message?

   Note: It is a system operator message if the display indicates that the message is in the QSYSOPR message queue. Critical messages can be found in the QSYSMSG message queue. For more information, see the *Create message queue QSYSMSG for severe messages* topic in the Troubleshooting navigation category of IBM i Knowledge Center.

   - **Yes:** Continue with the next step.
   - **No:** Go to step "8" on page 8.

---

7. Is the system operator message highlighted, or does it have an asterisk (*) next to it?

   - **Yes:** Go to step "17" on page 9.
   - **No:** Go to step "12" on page 9.

---

8. Move the cursor to the message line and press F1 (Help). Does the Additional Message Information display appear?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "10" on page 8.

---

9. Record the additional message information on the appropriate problem reporting form. For details, see "Problem reporting form" on page 22.

   Follow the recovery instructions on the Additional Message Information display.

   Did this solve the problem?

   - **Yes: This ends the procedure**.
   - **No:** Continue with the next step.

---

10. To display system operator messages, type `dspmsg qsysopr` on any command line and then press **Enter**.

    Did you find a message that is highlighted or has an asterisk (*) next to it?

    - **Yes:** Go to step "17" on page 9.
    - **No:** Continue with the next step.

    Note: The message monitor in the IBM Navigator for i Web console can also inform you when a problem has developed. For details, see the *Scenario: Message monitor topic in the Systems Management navigation category* of IBM i Knowledge Center.

---

11. Did you find a message with a date or time that is at or near the time the problem occurred?

Note: Move the cursor to the message line and press F1 (Help) to determine the time that a message occurred. If the problem is shown to affect only one console, you might be able to use information from the JOB menu to diagnose and solve the problem. To find this menu, type **GO JOB** and press **Enter** on any command line.

- **Yes:** Continue with the next step.
- **No:** Go to step "14" on page 9.

---

12. Complete the following steps:

    a. Move the cursor to the message line and press F1 (Help) to display additional information about the message.

    b. Record the additional message information on the appropriate problem reporting form. For details, see "Problem reporting form" on page 22.

    c. Follow any recovery instructions that are shown.

    Did this solve the problem?

    - **Yes: This ends the procedure**.
    - **No:** Continue with the next step.

---

13. Did the message information indicate to look for additional messages in the system operator message queue (QSYSOPR)?

    - **Yes:** Press F12 (Cancel) to return to the list of messages and look for other related messages. Then, return to step "10" on page 8.
    - **No:** Continue with the next step.

---

14. Do you know which input/output device is causing the problem?

    - **Yes:** Continue with step "16" on page 9.
    - **No:** Continue with the next step.

---

15. If you do not know which input/output device is causing the problem, describe the problems that you observed by completing the following steps:

    a. Type `GO USERHELP` on any command line and then press **Enter**.

    b. Select option 10 (Save information to help resolve a problem).

    c. Type a brief description of the problem and then press **Enter**. If you specify the default **Y** in the **Enter notes about problem** field, you can enter more text to describe your problem.

    d. Report the problem to your hardware service provider.

---

16. Complete the following steps:

    a. Type ANZPRB on the command line and then press **Enter**. For details, see *Using the Analyze Problem (ANZPRB) command* in the Troubleshooting navigation category in IBM i Knowledge Center.

    b. Contact your next level of support. **This ends the procedure**.

    Note: To describe your problem in greater detail, see *Using the Analyze Problem (ANZPRB) command* in the Troubleshooting navigation category in IBM i Knowledge Center. This command can also run a test to further isolate the problem.

---

17. Complete the following steps:

a. Move the cursor to the message line and press F1 (Help) to display additional information about the message.

b. Press F14, or use the Work with Problem (WRKPRB) command. For details, see *Work with Problem (WRKPRB)* in the Troubleshooting navigation category in IBM i Knowledge Center.

c. If this does not solve the problem, see the Symptom and recovery actions.

---

18. Choose from the following options:

- If reference codes appear on the control panel or the management console, record them. Then, go to the Reference code finder to see if additional details are available for the code you received.
- If no reference codes appear on the control panel or the management console, a serviceable event is indicated by a message in the problem log. Use the WRKPRB command. For details, see *Work with Problem (WRKPRB)* in the Troubleshooting navigation category in IBM i Knowledge Center.

---

19. Details about errors that occur when IBM i is not running or when IBM i is now not accessible can be found in the control panel or in the Advanced System Management Interface (ASMI).

Do you choose to look for error details by using ASMI?

- **Yes:** Go to step “21” on page 10.
- **No:** Continue with the next step.

---

20. At the control panel, complete the following steps.

a. Press the increment or decrement button until the number 11 is displayed in the upper-left corner of the display.

b. Press **Enter** to display the contents of function 11.

c. Look for a reference code in the upper-right corner.

Is there a reference code displayed on the control panel in function 11?

- **Yes:** Go to step “22” on page 10.
- **No:** Contact your hardware service provider. **This ends the procedure.**

---

21. On the console that is connected to the ASMI, complete the following steps.

Note: If you are unable to locate the reported problem, and there is more than one open problem near the time of the reported failure, use the earliest problem in the log.

a. Log in with a user ID that has an authority level as general, administrator, or authorized service provider.

b. In the navigation area, expand **System Service Aids** and click **Error/Event Logs**. If log entries exist, a list of error and event log entries is displayed in a summary view.

c. Scroll through the log under **Serviceable Customer Attention Events** and verify that there is a problem to correspond with the failure.

For more detailed information on the ASMI, see Managing the Advanced System Management Interface.

Do you find a serviceable event, or an open problem near the time of the failure?

- **Yes:** Continue with the next step.
- **No:** Contact your hardware service provider. **This ends the procedure.**

---

22. The reference code description might provide information or an action that you can take to correct the failure.

Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

For more information about reference codes, see Reference codes.

Was there a reference code description that helped you to resolve the problem?

- **Yes: This ends the procedure.**
- **No:** Continue with the next step.

23. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

    - If a field-replaceable unit (FRU) location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
    - If an isolation procedure is listed for the reference code in the reference code lookup information, include the isolation procedure as a corrective action even if it is not listed in the serviceable event view or control panel.
    - If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

    To find error details on the control panel, complete the following steps:

    a. Press **Enter** to display the contents of function 14. If data is available in function 14, the reference code has a FRU list.
    b. Record the information in functions 11 through 20 on the control panel.
    c. Contact your service provider and report the reference code and other information.

    To find error details on the ASMI, complete the following steps from the Error Event Log view:

    a. Record the reference code.
    b. Select the corresponding check box on the log and click Show details.
    c. Record the error details.
    d. Contact your service provider.

    This ends the procedure.

Light path diagnostics on Power Systems

Light path diagnostics is a simplified approach for repair actions on Power Systems hardware that provides fault indicators to identify parts that need to be replaced.

Light path diagnostics is a system of light-emitting diodes (LEDs) on the control panel and on various internal components of the Power Systems hardware. When an error occurs, LEDs are lit throughout the system to help identify the source of the error.

With light path diagnostics, the fault LED for the FRUs to be replaced will be active when the unit is powered on. The failing FRUs can be attached to another FRU, which must be first removed to access the failing FRUs. For those cases, light path diagnostics provides a blue switch on the FRU that has to be removed first. When the first FRU is removed, you can press and hold the light path diagnostics switch to light the LEDs and locate the failing part. In most of the situations, the switch will have enough power stored to activate the LEDs for two hours after the unit has been powered off. However, this can vary significantly and therefore the switch should be used as soon as possible. The amber LEDs can normally be kept active for 30 seconds, however, this can also vary. Associated with the light path diagnostics switch is a green LED that will be activated when the switch is used and there is enough power stored to activate the amber LEDs. If the green LED does not activate when the switch is pressed then there is not enough power remaining to activate any amber LEDs on that FRU. If that happens then light path diagnostics and FRU identify function cannot be used for replacing the failing FRUs. Perform the repair action using the location codes in the error log or if determined by problem analysis as if the unit did not have light path diagnostics and did not have functioning identify LEDs.

At the core of light path diagnostics is a set of fault indicators that are implemented as amber LEDs. These LEDs provide a way for the diagnostics to identify which field-replaceable unit (FRU) needs to be replaced. Service labels, color-coded touch points for the FRUs, and a no tools required design for FRU removal and installation are all elements of light path diagnostics.

With light path diagnostics, at the same time that the diagnostics create an error log for the problem, it also activates the fault indicator when a FRU has a failed or failing component. This includes predictive failure analysis (PFA). The FRU fault LED is turned on solid (not flashing). Whenever a fault indicator is activated the enclosure's external Fault indicator on the operator panel is also turned on solid. The enclosure fault indicator on the panel means that inside the unit one or more FRU fault indicators is on. The error log shows the part number and location code of the FRU that must be replaced along with other FRUs or procedures to follow if replacing the first FRU does not resolve the problem.

If the diagnostics determine that the problem is firmware related, configuration related, or not isolated to a specific FRU then no fault indicator is activated. For these kinds of problems the amber System Info indicator on the operator panel is activated. The error log shows the procedures to follow and the possible FRUs that could be causing the problem.

During the repair action, the service label on the service access cover shows the FRUs and the steps required to remove or install the FRUs. Therefore, the basic flow of the repair is that the LEDs show which part to replace, the color-coded touch points indicate if the unit must be powered off to remove or install the part, and the service label shows the steps needed on the touch points. The FRU fault LED is not an indication that the FRU is ready to be replaced. To replace the FRU, some preparation steps might be necessary, such as removing the resource from use or powering off the unit. The service label and the touch point colors give the initial guidance for FRU removal.

When a FRU has been replaced, the fault indicator automatically turns off either when the new FRU is installed, or when the power is restored to the new FRU. This automatic shut off might take several seconds to a minute as the new FRU is powered on, brought online and tested by the system. When there are no more FRU fault indicators on in an enclosure then the enclosure's Fault indicator on the operator panel turns off automatically.

In addition to the fault indicators, there are also amber identify indicators for each FRU. The identify indicators flash when activated. Identify indicators are used to help a servicer identify where a location is. The location may be occupied or empty. The servicer can turn them on and off from a user interface either during a repair action or during the installation of new parts or when removing parts. The identify indicator visually confirms where a location code is. Whenever an identify indicator is activated, the enclosure's blue *locate* or *beacon* LED on the operator panel is also activated (flashing).

The same amber LED on a FRU can be used for both fault and identify indications. Whenever the LED is on solid for a fault, the LED switches to flashing when the FRU identify function is turned on. When identify function is turned off, the LED returns to fault (solid on) if that was the previous state of the LED.

Replacing FRUs by using enclosure fault indicators

After you obtain a replacement part, use this procedure to identify the location of the part that needs to be replaced.

About this task

To identify and locate the part that requires replacement, complete the following steps.

Procedure

1. Before you move the unit into the service position, refer to the service label. It might be necessary to identify and remove cables that are attached to the FRU you are about to exchange. Go to “Service labels” on page 14 to locate the service label for your system. Use the FRU location code and the service label to determine whether there are any actions before you move the unit into the service position. Complete those actions and return to the next step in this procedure.

2. Identify the unit with the active enclosure fault indicator. Use the service label that is affixed to the service access cover and the amber light-emitting diode (LED) on the FRU to find the failing FRU. Move the unit into the service position, but do not remove the service access cover.

   - If the unit is rack-mounted, the service label is visible on the service access cover. Continue with the next step.
   - If the unit is a stand-alone system, the exterior cover must be removed to view the service label. Remove the exterior cover by using the procedure found in the following table and then return to the next step in this procedure.

   Table 1. Exterior cover removal procedures for the stand-alone servers

Machine type and model	Removal procedure
9009-41A	Removing the service access cover from a stand-alone 9009-41A system.

3. Using the service label, determine whether the FRU you are replacing can be exchanged without removing the service access cover. Is the FRU fault LED visible externally and active (on solid, not flashing) and does the service label show that the service access cover does not need to be removed to exchange the FRU? (Choose No if you are unsure.)

   Note: If you used the identify function in a user interface to help locate the FRU, then the amber LED is flashing. Otherwise, the amber LED is solid (not flashing).

   - **Yes:** Go to step “6” on page 13.
   - **No:** To identify which FRU to exchange, you must remove the service access cover and locate the FRU that has an active FRU fault indicator (amber LED on). Continue with the next step.

4. Remove the service access cover and locate the FRU that has an active fault indicator (amber LED on, not flashing). Use the following table to determine whether you must power off the unit before you remove the cover.

   Note: You can remove the service access cover while the unit is powered on.

5. Search for the FRU to be replaced by locating the active amber LED.

   Notes:

   - If you used the identify function in a user interface to help locate the FRU, then the amber LED is flashing. Otherwise, the amber LED is solid (not flashing).
   - Some FRUs might be an integral part of another FRU. This might make it difficult to see the FRU that you need to exchange or the amber LED that designates the FRU that needs to be exchanged. If so, remove all FRUs that are associated with the failing FRU.

   Do you need to remove another FRU to replace the FRU designated by the amber LED?

   - **Yes:** Go to step “9” on page 14.
   - **No:** Continue with the next step.

6. For the FRU you located with its fault LED active, was it replaced for this problem or service action?

   - **Yes:** The FRU replaced for the original problem did not resolve the problem. Go back to the serviceable event for the original problem and address the remaining FRUs that are listed.

     Note: If the fault indicator for the replaced FRU is turned on, use the Advanced System Management Interface (ASMI) to turn off the fault indicator.

     This ends the procedure.

   - **No:** Continue with the next step.

7. For the FRU you located with the active fault LED, compare the location code that you recorded for the replacement FRU of the problem you are working on with the location code of the active fault indicator. If they do not match, you are working a problem from the log that is different from the problem that activated the fault indicators. Do the location codes match?

- **Yes:** You are working the same problem that activated the fault indicators. Continue with the next step.
- **No:** If you have the correct replacement FRU for where the fault indicator is active, you can continue with this repair action. When you replace the FRU, record the location codes of the active fault indicators for use later to identify which problem to close when the repair is complete, then continue with the next step. Otherwise, contact your service provider to obtain the replacement part for the FRU with an active fault indicator and begin problem analysis again. **This ends the procedure.**

8. If you have not already done so, record the location of the FRU you are about to exchange either by where the service label shows the FRU or where it is in the unit by the location labeling. For information on part locations and location codes, see Part locations and location codes. In the locations and addresses information for your system, locate the FRU and the corresponding FRU exchange procedure. The exchange procedure provides the steps necessary to exchange the FRU. If the FRU can be replaced while the unit is powered on, the exchange procedure provides that option and the necessary instructions. If the enclosure fault indicator turns on again within a few minutes of completing the replacement and returning to normal use of the unit, then begin problem analysis again. Otherwise, close the problem. **This ends the procedure.**
9. If you have not already done so, record the location of the FRU that you plan to remove either by where the service label shows the FRU or by the location labeling where it is in the unit. For information on part locations and location codes, see Part locations and location codes.

   In the locations and addresses information, locate the FRU and the corresponding FRU exchange procedure. The exchange procedure provides the steps necessary to remove this FRU. If the FRU can be removed while the unit is powered on, the exchange procedure provides that option and the necessary instructions. When you remove this FRU, the fault indicator for the associated FRU you are exchanging turns off. This FRU has an LED activation button that you can press that powers the amber indicator of the FRU you are exchanging. Use the button to locate the FRU you are exchanging. Go to step “8” on page 14.

   Note: If the button's green LED does not activate, there is not enough charge in the switch to activate the amber fault LED. To identify the failing FRU, use the FRU location code either from the error log or as determined by problem analysis.
10. For the FRU you located that has its fault LED active, were any of them replaced for this problem or service action?

    - **Yes:** The FRU replaced for the original problem did not resolve the problem. Go back to the serviceable event for the original problem and work the remaining FRUs listed. Use ASMI to turn off the fault indicator for the FRU. **This ends the procedure.**
    - **No:** Use the information on the service label to exchange the FRU. When the FRU is exchanged, use the service label to guide you in reassembling the unit. Power on the unit. The FRU fault indicator is turned off during the power-on process if it was not already turned off. If the enclosure fault indicator turns on again within a few minutes of powering on the unit, then begin problem analysis again. Otherwise, close this problem. **This ends the procedure.**

Service labels

Use this information to view the service labels on system models or expansion units.

Service label for the 5105-22E

Service labels identify service locations on system models or expansion units.

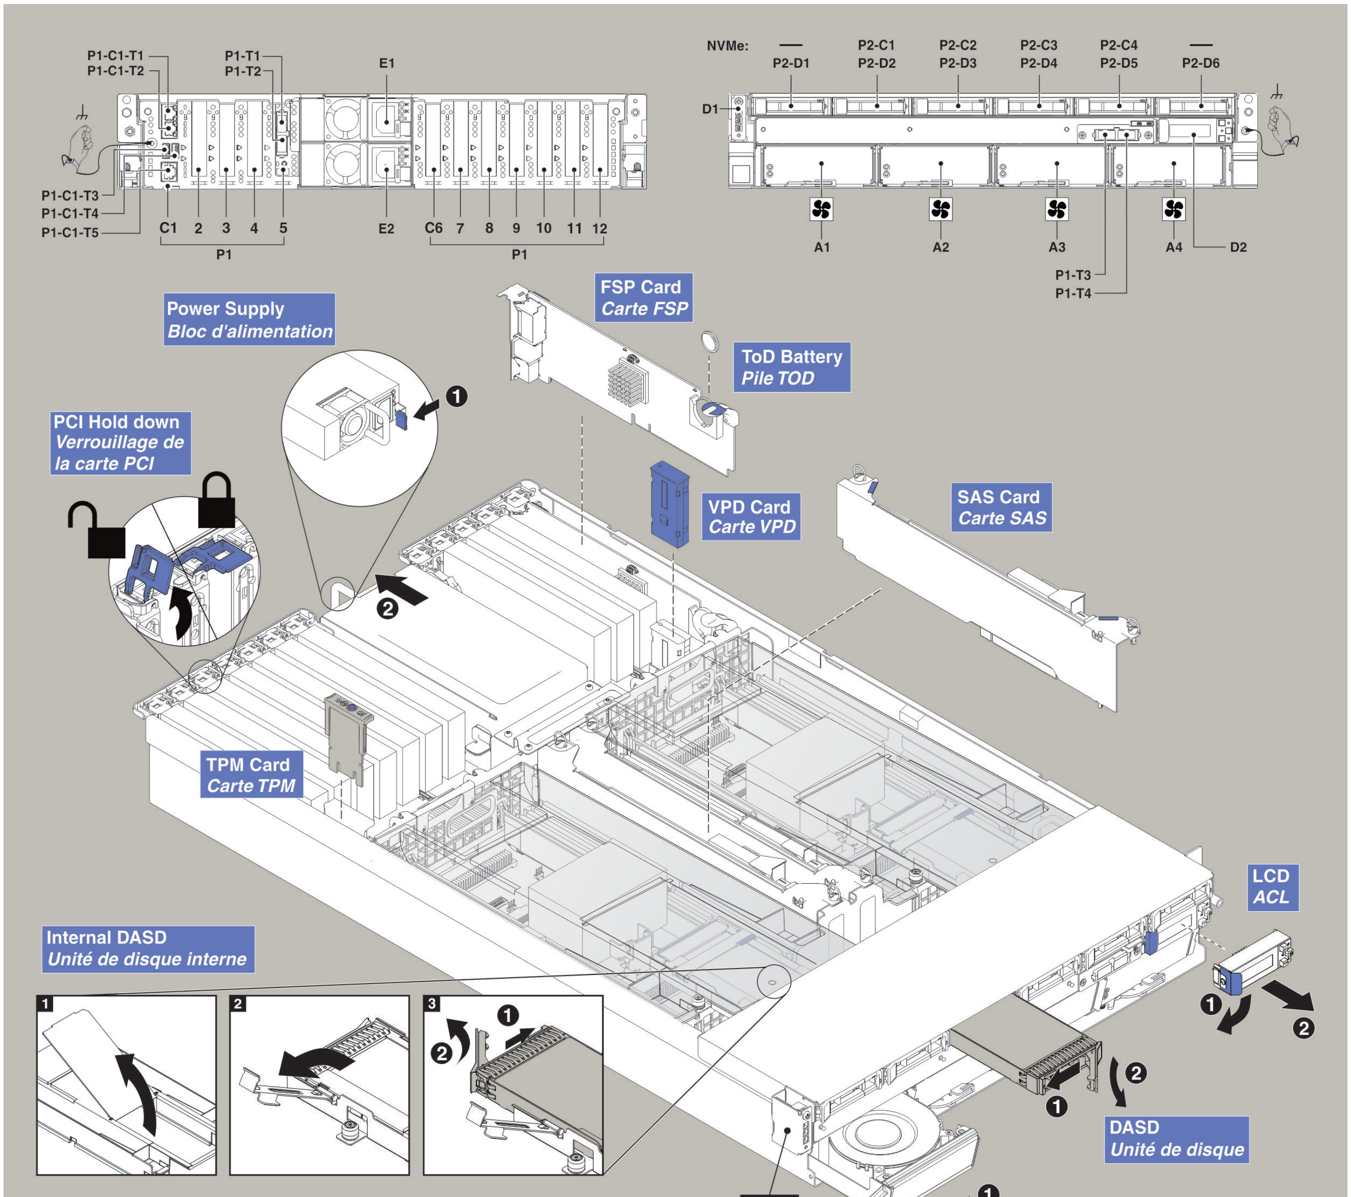

⏻	Power button - From standby mode, push for power ON - If power ON, push and hold for standby mode. *Bouton d'alimentation* *- À partir du mode de veille, appuyer pour mettre sous tension* *- Si sous tension, appuyer et maintenir pour mettre en mode de veille*
	System locate LED - Blue indicates identify *Voyant DEL de localisation du système – un témoin bleu indique l'état d'identification.*
	Check log LED - Solid amber indicates check for entry in error log *Voyant DEL de vérification du journal – un témoin ambre continu indique de vérifier l'entrée dans le journal d'erreurs.*
!	System error LED - Solid amber indicates enclosure fault, attention required *Voyant DEL d'erreur système – un témoin ambre continu indique une défaillance qui doit être corrigée.*

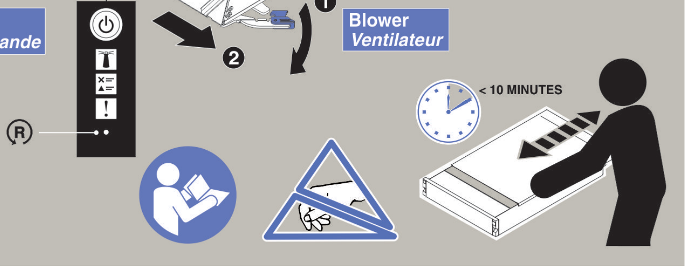

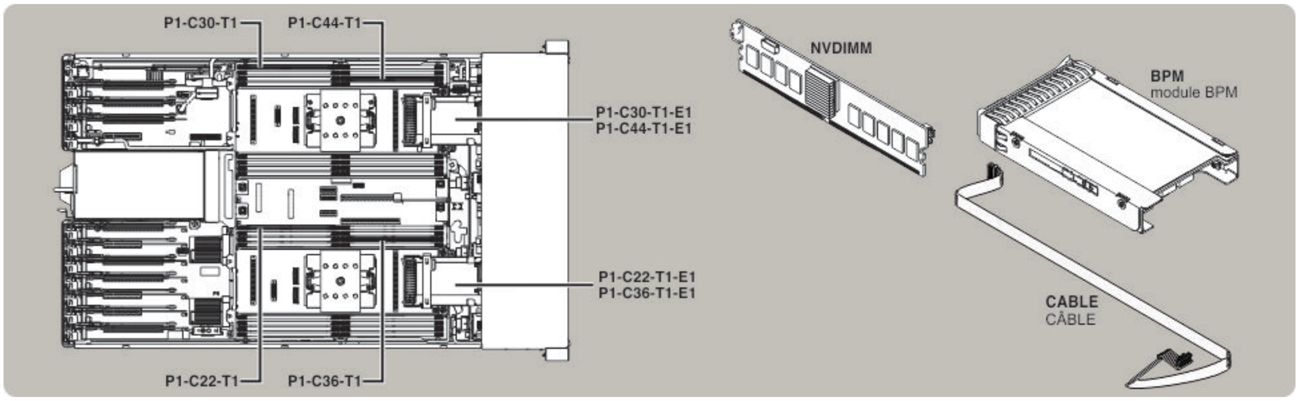
P1-C30-T1
P1-C44-T1
P1-C30-T1-E1
P1-C44-T1-E1
P1-C22-T1-E1
P1-C36-T1-E1
P1-C22-T1
P1-C36-T1
NVDIMM
BPM
module BPM
CABLE
CÂBLE

Service label for the 9009-22G or 9223-22S

Service labels identify service locations on system models or expansion units.

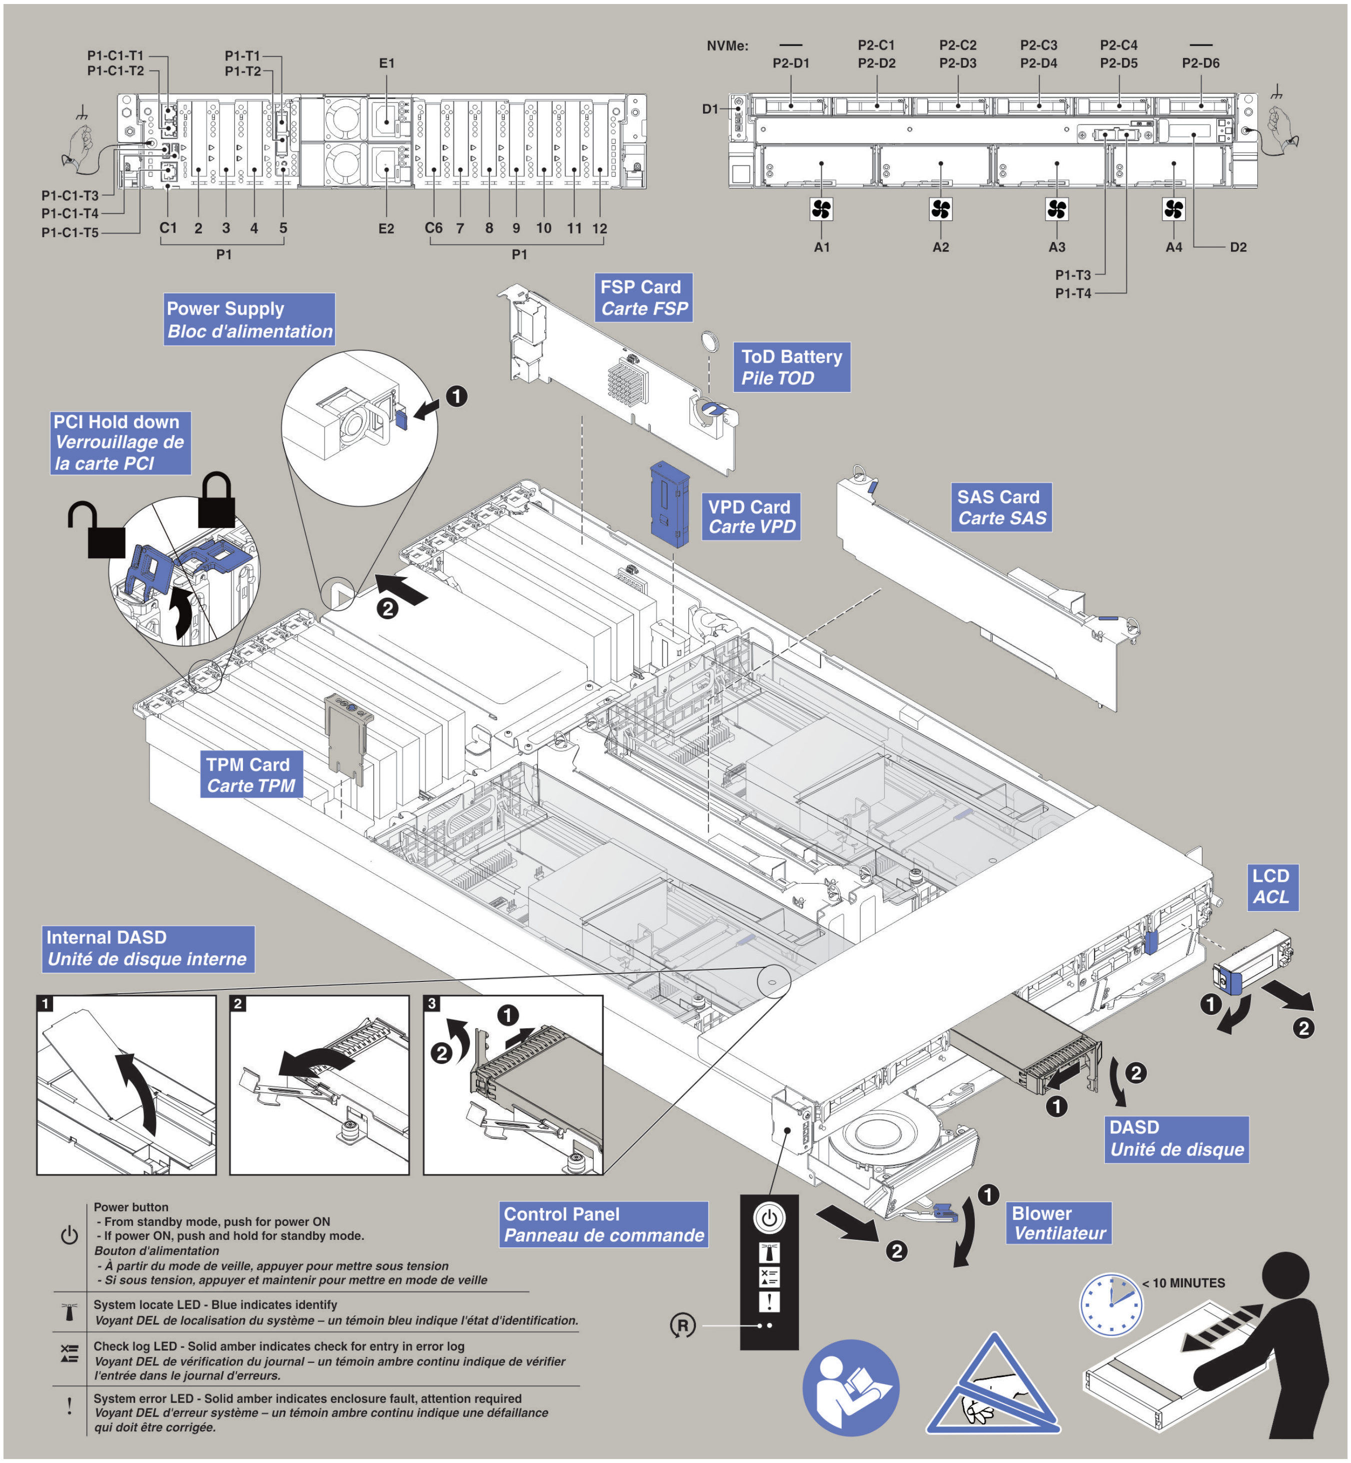

⏻ Power button
- From standby mode, push for power ON
- If power ON, push and hold for standby mode.

Bouton d'alimentation
- *À partir du mode de veille, appuyer pour mettre sous tension*
- *Si sous tension, appuyer et maintenir pour mettre en mode de veille*

System locate LED - Blue indicates identify
Voyant DEL de localisation du système – un témoin bleu indique l'état d'identification.

Check log LED - Solid amber indicates check for entry in error log
Voyant DEL de vérification du journal – un témoin ambre continu indique de vérifier l'entrée dans le journal d'erreurs.

! System error LED - Solid amber indicates enclosure fault, attention required
Voyant DEL d'erreur système – un témoin ambre continu indique une défaillance qui doit être corrigée.

Service label for the 9008-22L, 9009-22A, or 9223-22H

Service labels identify service locations on system models or expansion units.

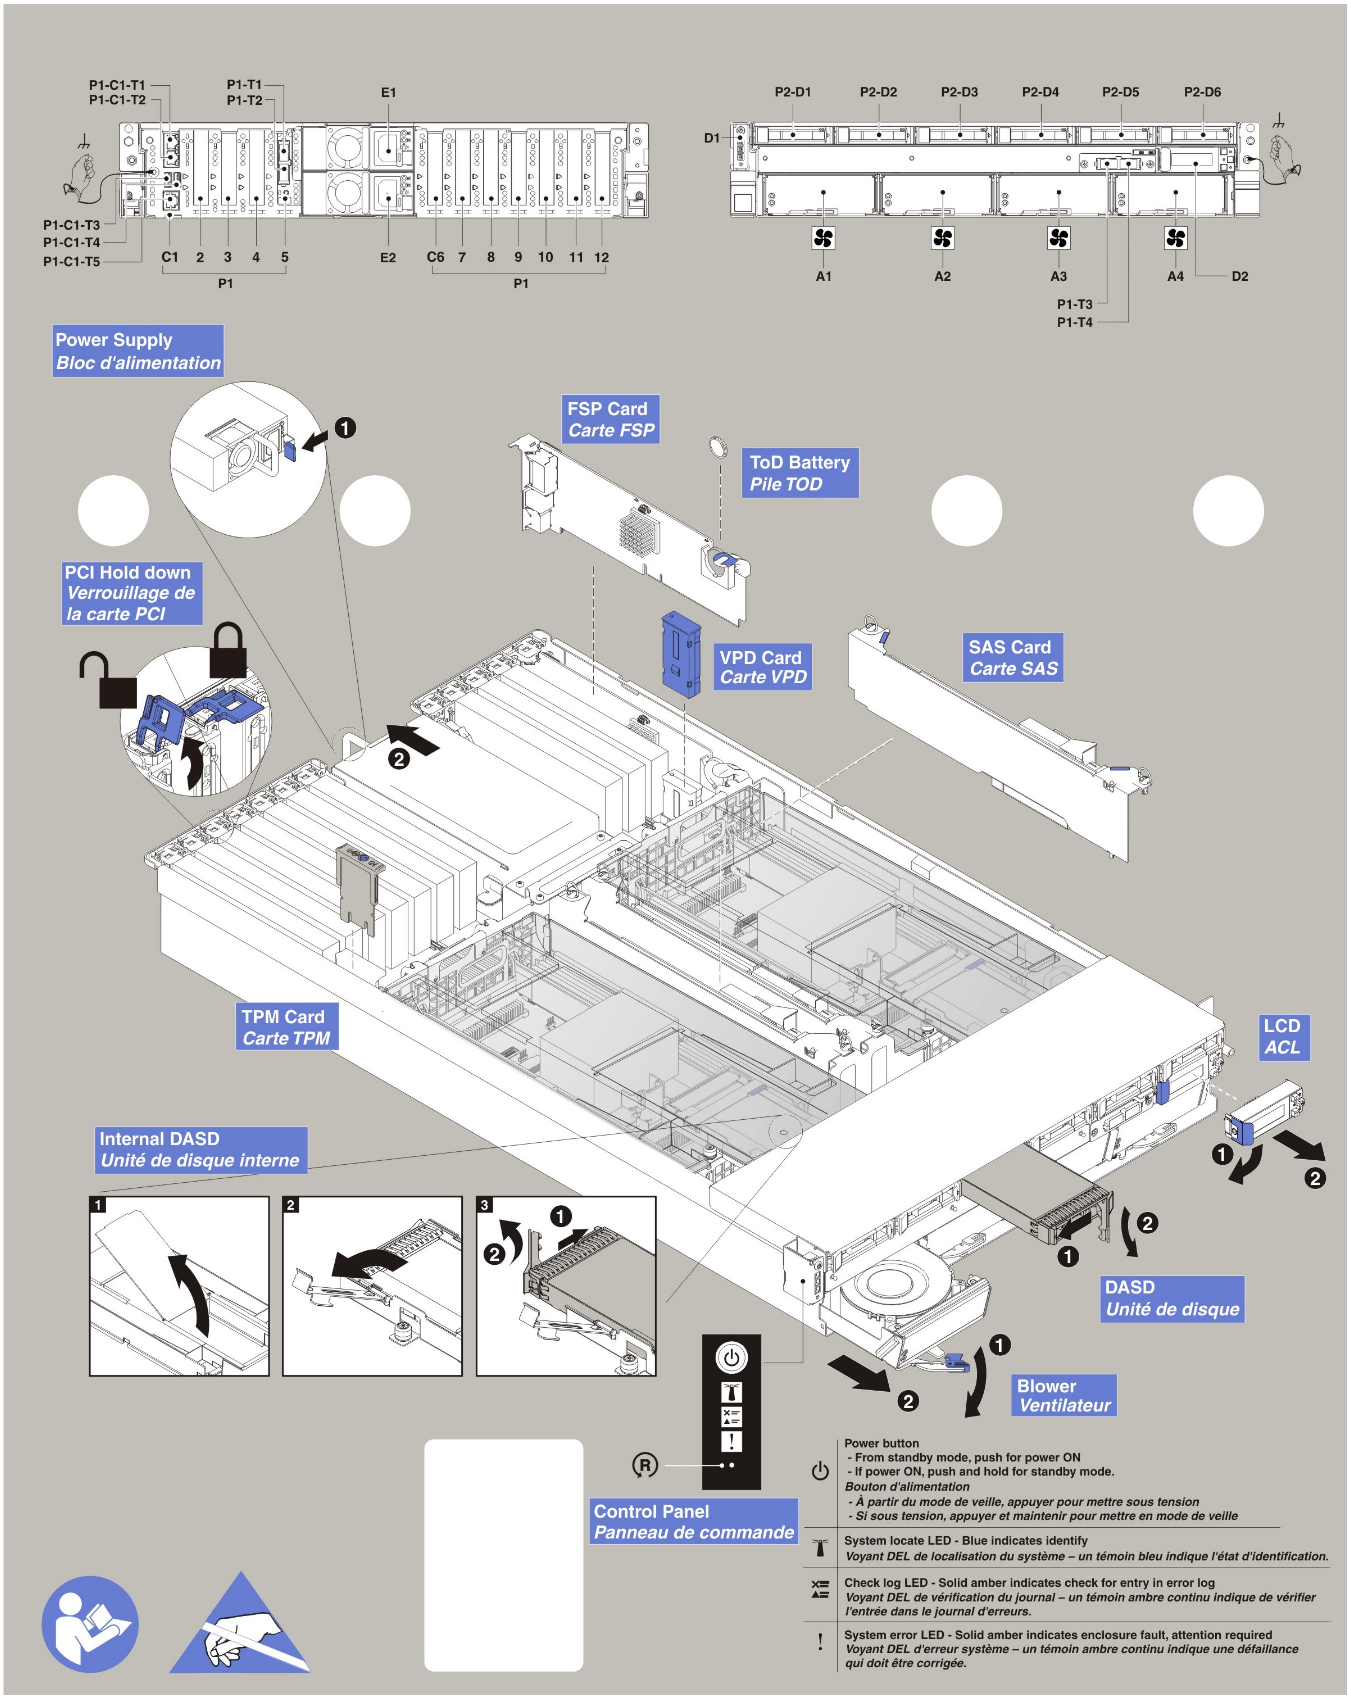

Power button
- From standby mode, push for power ON
- If power ON, push and hold for standby mode.

Bouton d'alimentation
- *À partir du mode de veille, appuyer pour mettre sous tension*
- *Si sous tension, appuyer et maintenir pour mettre en mode de veille*

System locate LED - Blue indicates identify
Voyant DEL de localisation du système – un témoin bleu indique l'état d'identification.

Check log LED - Solid amber indicates check for entry in error log
Voyant DEL de vérification du journal – un témoin ambre continu indique de vérifier l'entrée dans le journal d'erreurs.

System error LED - Solid amber indicates enclosure fault, attention required
Voyant DEL d'erreur système – un témoin ambre continu indique une défaillance qui doit être corrigée.

Service label for the 9009-41A, 9009-42A, or 9223-42H

Service labels identify service locations on system models or expansion units.

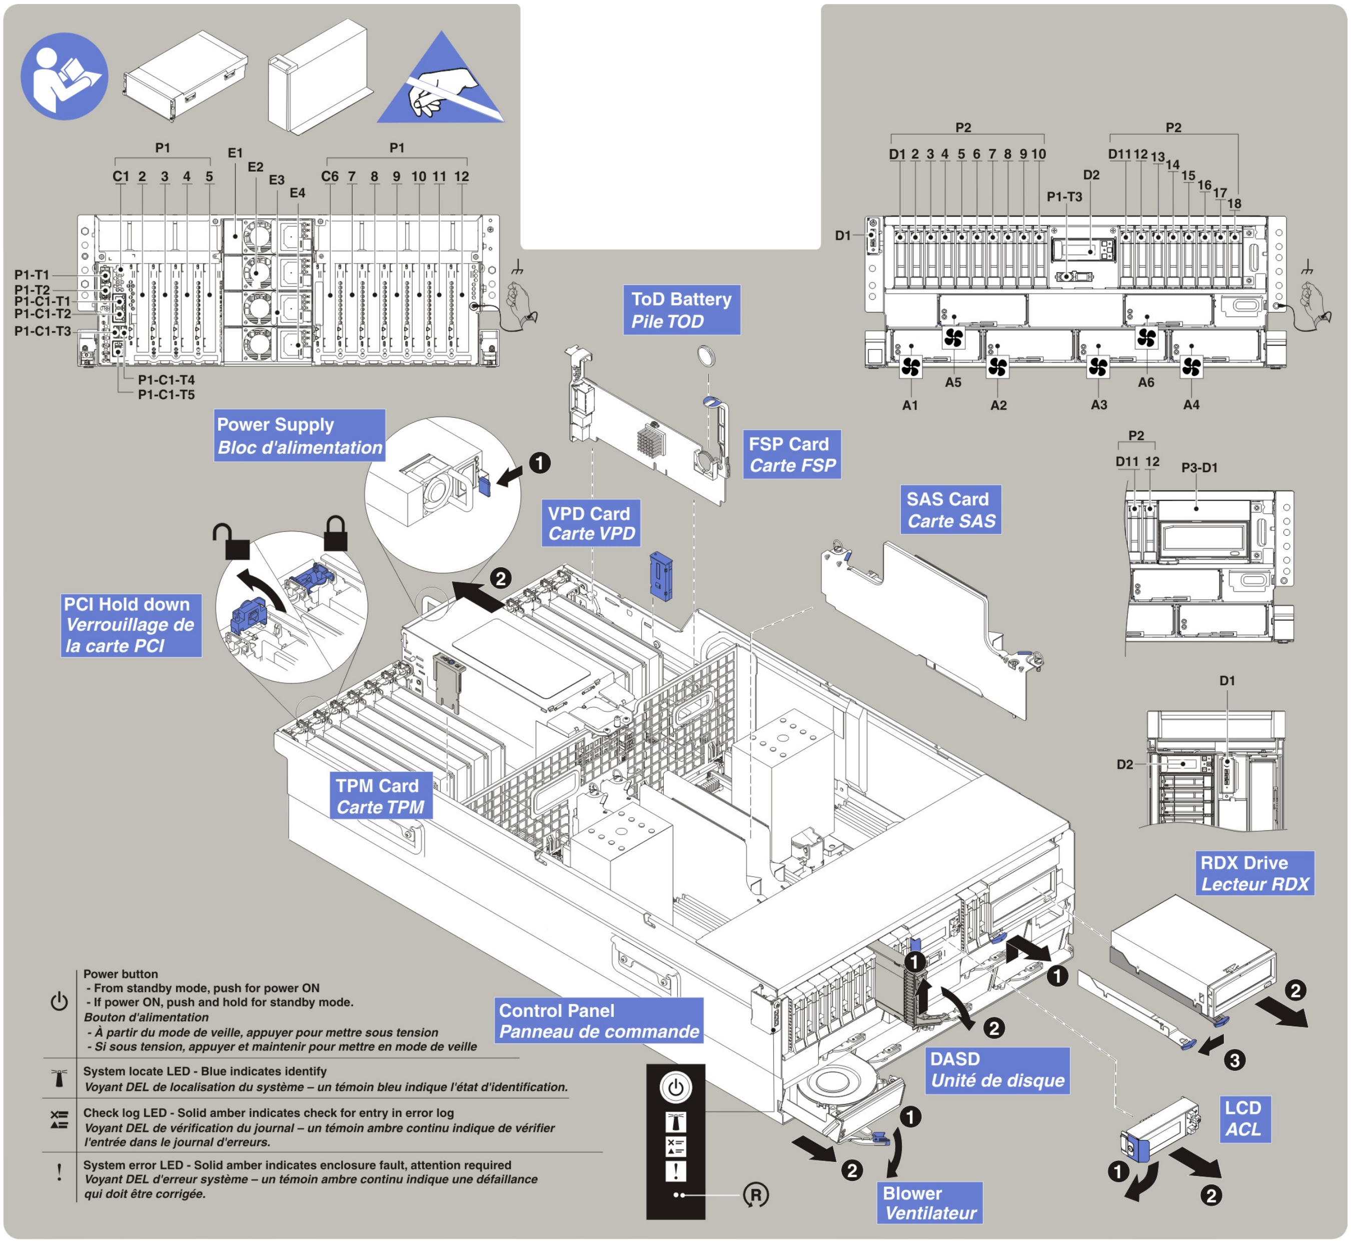

⏻	Power button - From standby mode, push for power ON - If power ON, push and hold for standby mode. *Bouton d'alimentation* *- À partir du mode de veille, appuyer pour mettre sous tension* *- Si sous tension, appuyer et maintenir pour mettre en mode de veille*
	System locate LED - Blue indicates identify *Voyant DEL de localisation du système – un témoin bleu indique l'état d'identification.*
	Check log LED - Solid amber indicates check for entry in error log *Voyant DEL de vérification du journal – un témoin ambre continu indique de vérifier l'entrée dans le journal d'erreurs.*
!	System error LED - Solid amber indicates enclosure fault, attention required *Voyant DEL d'erreur système – un témoin ambre continu indique une défaillance qui doit être corrigée.*

Service label for the 9009-41G, 9009-42G, or 9223-42S

Service labels identify service locations on system models or expansion units.

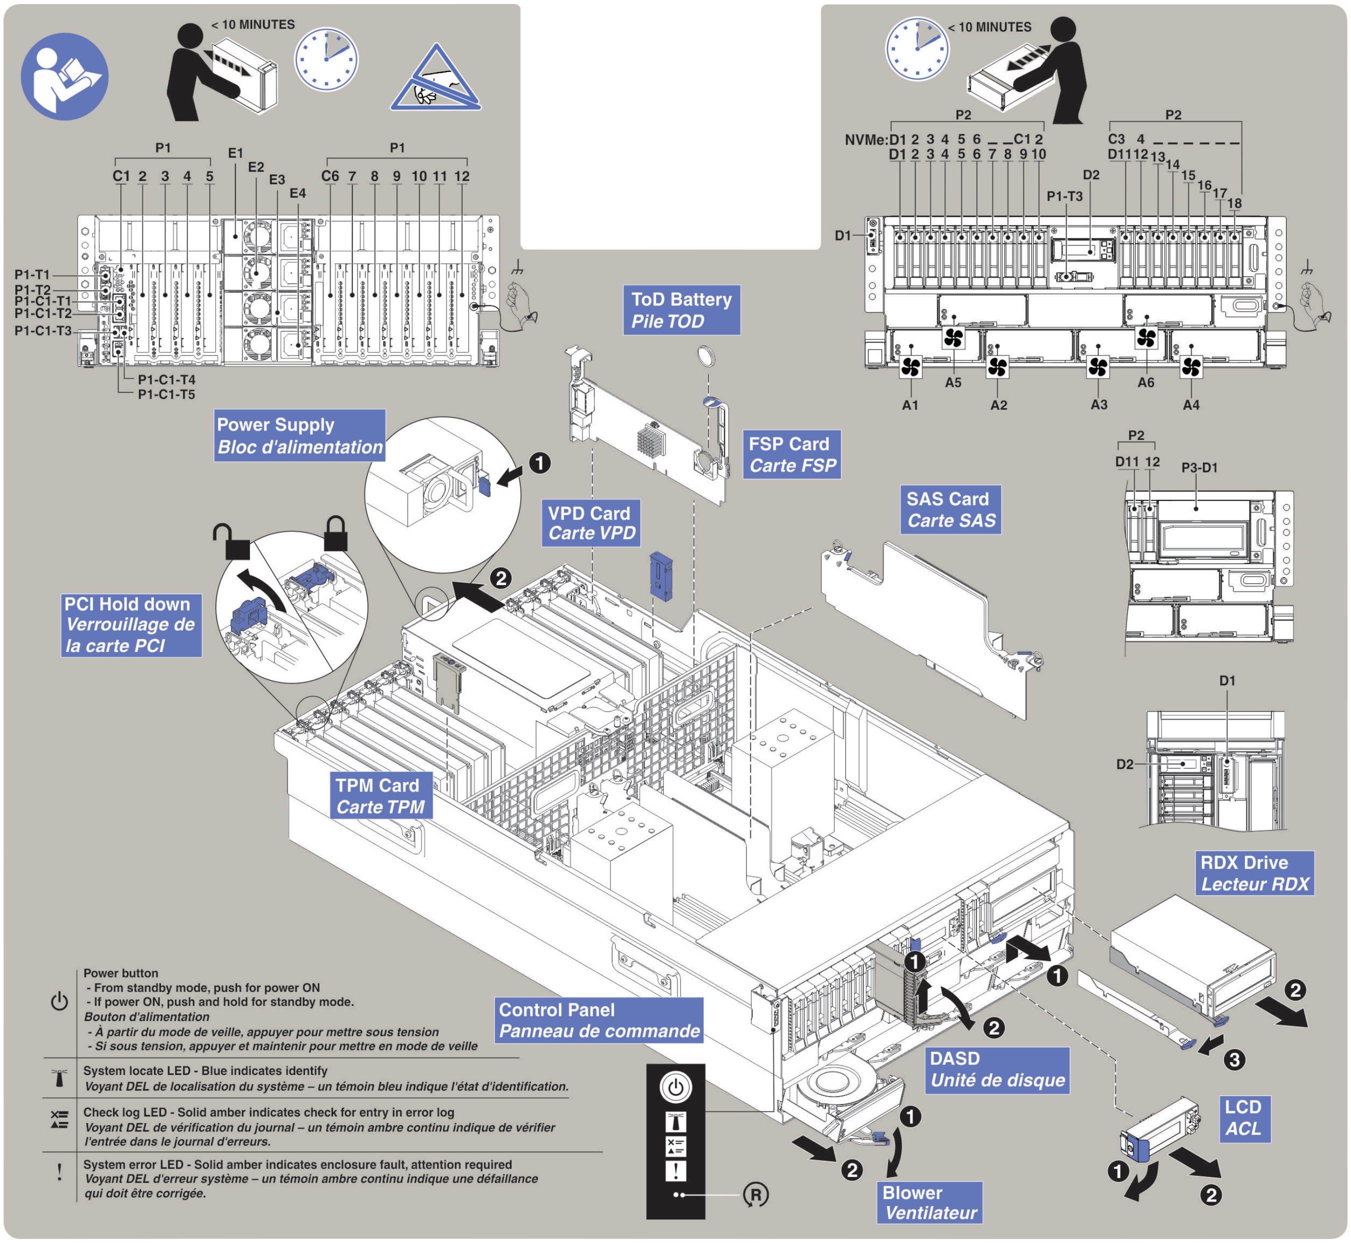

Power button
- From standby mode, push for power ON
- If power ON, push and hold for standby mode.

Bouton d'alimentation
- *À partir du mode de veille, appuyer pour mettre sous tension*
- *Si sous tension, appuyer et maintenir pour mettre en mode de veille*

System locate LED - Blue indicates identify
Voyant DEL de localisation du système – un témoin bleu indique l'état d'identification.

Check log LED - Solid amber indicates check for entry in error log
Voyant DEL de vérification du journal – un témoin ambre continu indique de vérifier l'entrée dans le journal d'erreurs.

System error LED - Solid amber indicates enclosure fault, attention required
Voyant DEL d'erreur système – un témoin ambre continu indique une défaillance qui doit être corrigée.

Service label for the 9040-MR9

Service labels identify service locations on system models or expansion units.

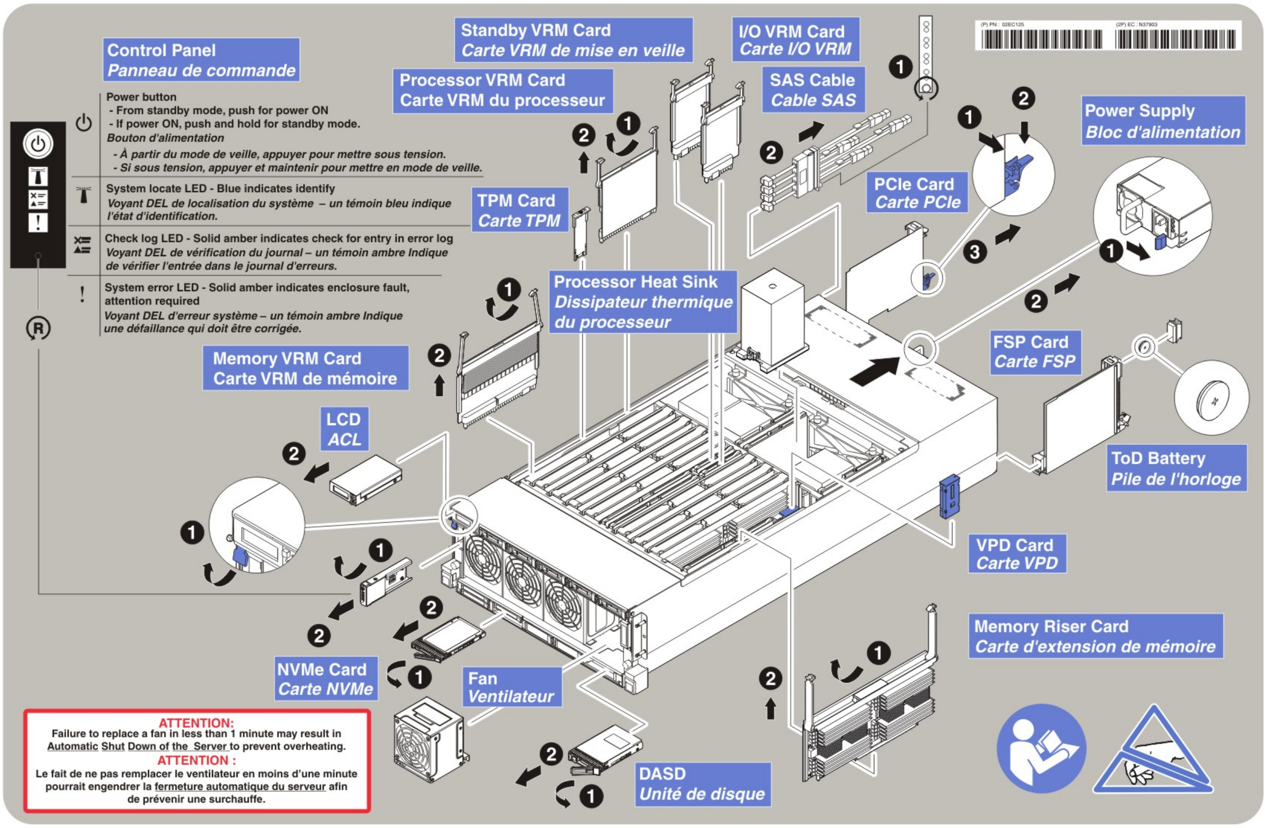

Power button
- From standby mode, push for power ON
- If power ON, push and hold for standby mode.

Bouton d'alimentation
- *À partir du mode de veille, appuyer pour mettre sous tension.*
- *Si sous tension, appuyer et maintenir pour mettre en mode de veille.*

System locate LED - Blue indicates identify
Voyant DEL de localisation du système – un témoin bleu indique l'état d'identification.

Check log LED - Solid amber indicates check for entry in error log
Voyant DEL de vérification du journal – un témoin ambre Indique de vérifier l'entrée dans le journal d'erreurs.

System error LED - Solid amber indicates enclosure fault, attention required
Voyant DEL d'erreur système – un témoin ambre Indique une défaillance qui doit être corrigée.

ATTENTION:
Failure to replace a fan in less than 1 minute may result in Automatic Shut Down of the Server to prevent overheating.
ATTENTION :
Le fait de ne pas remplacer le ventilateur en moins d'une minute pourrait engendrer la fermeture automatique du serveur afin de prévenir une surchauffe.

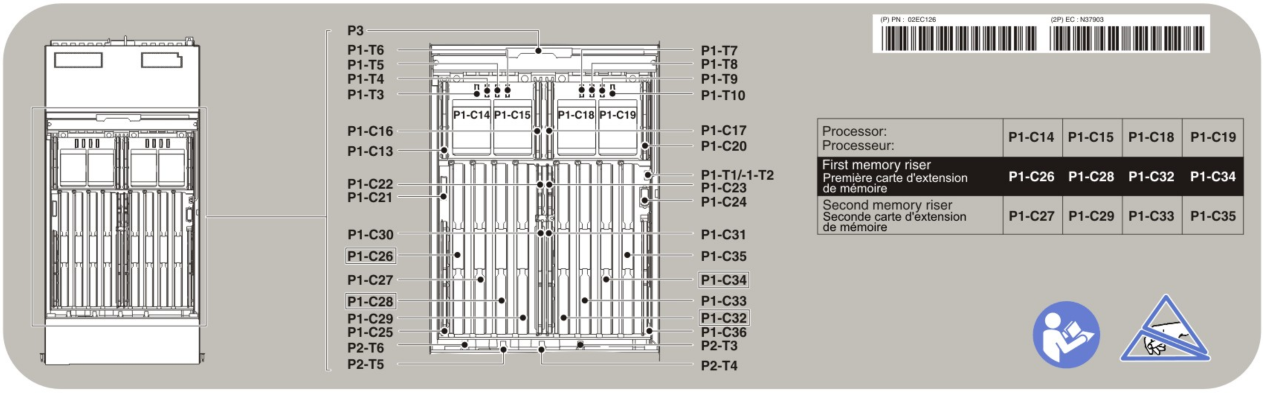

Processor: Processeur:	P1-C14	P1-C15	P1-C18	P1-C19
First memory riser Première carte d'extension de mémoire	P1-C26	P1-C28	P1-C32	P1-C34
Second memory riser Seconde carte d'extension de mémoire	P1-C27	P1-C29	P1-C33	P1-C35

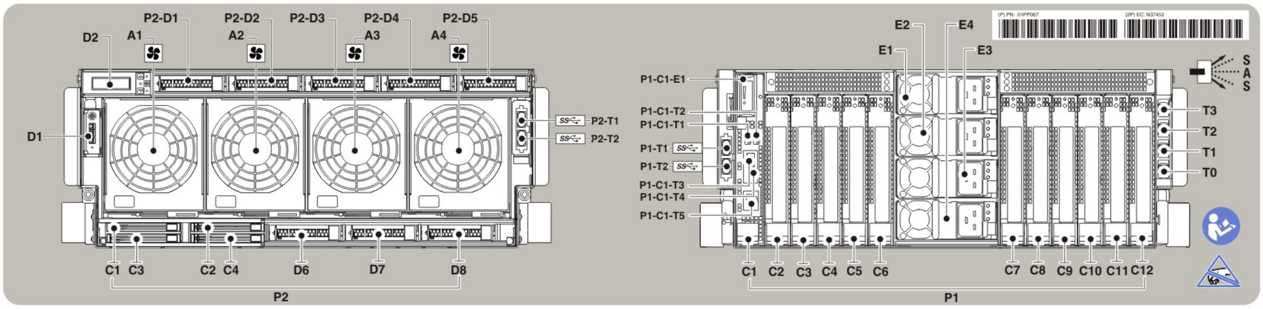

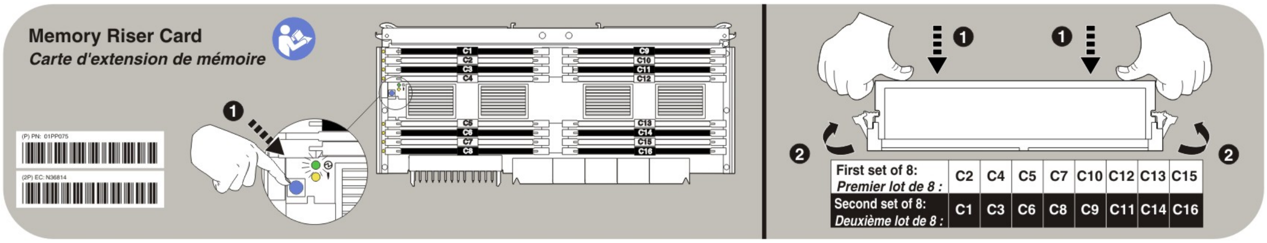

Memory Riser Card

Carte d'extension de mémoire

(P) PN: 01PP075

(2P) EC: N36814

First set of 8: *Premier lot de 8 :*	C2	C4	C5	C7	C10	C12	C13	C15
Second set of 8: *Deuxième lot de 8 :*	C1	C3	C6	C8	C9	C11	C14	C16

Service label for the EMX0 PCIe Gen3 I/O expansion drawer

Service labels identify service locations on system models or expansion units.

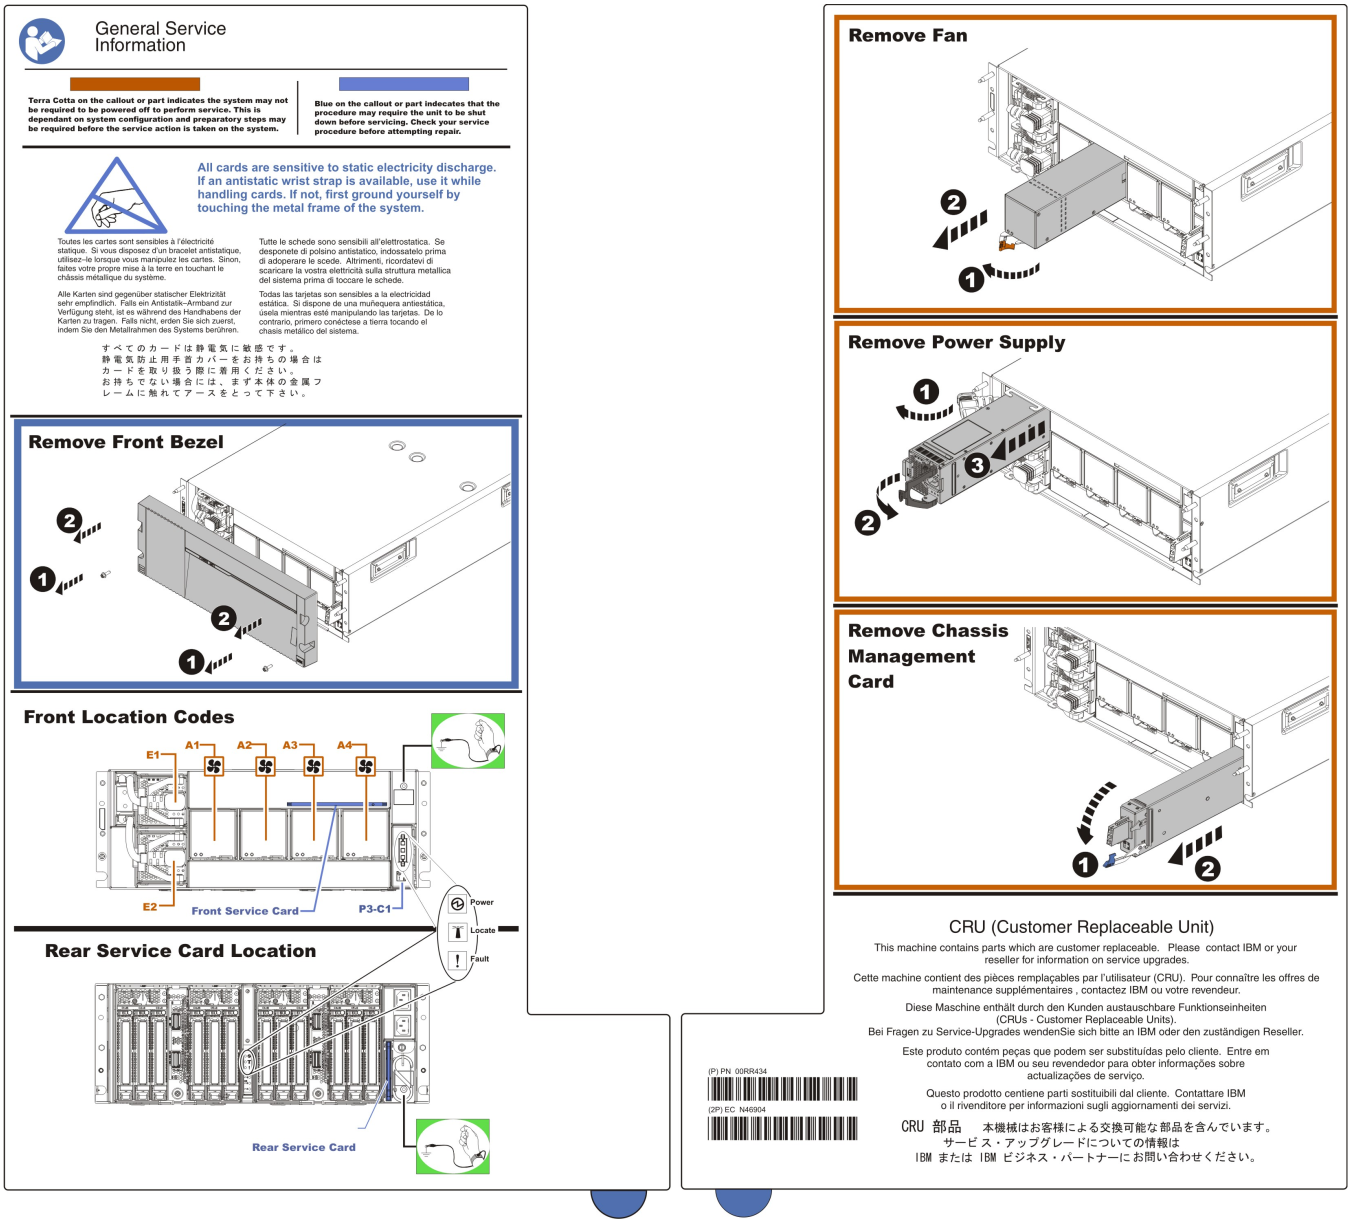

General Service Information

Terra Cotta on the callout or part indicates the system may not be required to be powered off to perform service. This is dependant on system configuration and preparatory steps may be required before the service action is taken on the system.

Blue on the callout or part indecates that the procedure may require the unit to be shut down before servicing. Check your service procedure before attempting repair.

All cards are sensitive to static electricity discharge. If an antistatic wrist strap is available, use it while handling cards. If not, first ground yourself by touching the metal frame of the system.

Toutes les cartes sont sensibles à l'électricité statique. Si vous disposez d'un bracelet antistatique, utilisez–le lorsque vous manipulez les cartes. Sinon, faites votre propre mise à la terre en touchant le châssis métallique du système.

Alle Karten sind gegenüber statischer Elektrizität sehr empfindlich. Falls ein Antistatik–Armband zur Verfügung steht, ist es während des Handhabens der Karten zu tragen. Falls nicht, erden Sie sich zuerst, indem Sie den Metallrahmen des Systems berühren.

Tutte le schede sono sensibili all'elettrostatica. Se desponete di polsino antistatico, indossatelo prima di adoperare le scede. Altrimenti, ricordatevi di scaricare la vostra elettricità sulla struttura metallica del sistema prima di toccare le schede.

Todas las tarjetas son sensibles a la electricidad estática. Si dispone de una muñequera antiestática, úsela mientras esté manipulando las tarjetas. De lo contrario, primero conéctese a tierra tocando el chasis metálico del sistema.

すべてのカードは静電気に敏感です。
静電気防止用手首カバーをお持ちの場合は
カードを取り扱う際に着用ください。
お持ちでない場合には、まず本体の金属フ
レームに触れてアースをとって下さい。

Remove Front Bezel

Front Location Codes

Rear Service Card Location

Remove Fan

Remove Power Supply

Remove Chassis Management Card

CRU (Customer Replaceable Unit)

This machine contains parts which are customer replaceable. Please contact IBM or your reseller for information on service upgrades.

Cette machine contient des pièces remplaçables par l'utilisateur (CRU). Pour connaître les offres de maintenance supplémentaires , contactez IBM ou votre revendeur.

Diese Maschine enthält durch den Kunden austauschbare Funktionseinheiten (CRUs - Customer Replaceable Units).
Bei Fragen zu Service-Upgrades wendenSie sich bitte an IBM oder den zuständigen Reseller.

Este produto contém peças que podem ser substituídas pelo cliente. Entre em contato com a IBM ou seu revendedor para obter informações sobre actualizações de serviço.

Questo prodotto centiene parti sostituibili dal cliente. Contattare IBM o il rivenditore per informazioni sugli aggiornamenti dei servizi.

CRU 部品　本機械はお客様による交換可能な部品を含んでいます。
サービス・アップグレードについての情報は
IBM または IBM ビジネス・パートナーにお問い合わせください。

(P) PN 00RR434

(2P) EC N46904

Problem reporting form

Use the problem reporting form to record information about your server that will assist you in problem analysis.

Collect as much information as possible in the tables below, using either the control panel or the management console to gather the information.

Table 2. Customer, system, and problem information

Customer information and problem description	
Your name	
Telephone number	
IBM customer number, if available	
Date and time that the problem occurred	
Describe the problem	
System Information	
Machine type	
Model	
Serial number	
IPL type	
IPL mode	
Message information	
Message ID	
Message text	
Service request number (SRN)	
In which mode were IBM hardware diagnostics run?	___ Online or ___ stand-alone ? ___ Service mode or ___ concurrent mode?

Go to the management console or the control panel and indicate whether the following lights are on.

Table 3. Control panel lights

Control panel light	**Place a check if light is on**
Power On	
System Attention	

Go to the management console or control panel to find and record the values for functions 11 through 20. Use the following grid to record the characters shown on the management console or Function/Data display.

Table 4. Function values

Function	**Value**
11	__ __
12	__ __
13	__ __

Table 4. Function values (continued)

Function	Value
14	__ __
15	__ __
16	__ __
17	__ __
18	__ __
19	__ __
20 (for control panel users)	__ __
20 (for management console users)	Machine type: Model: Processor feature code: IPL type:

Starting a repair action

This is the starting point for repair actions. All repair actions must begin with this procedure. From this point, you are guided to the appropriate information to help you perform the necessary steps to repair the server.

Note: In this topic, **control panel** and **operator panel** are synonymous.

Before beginning, record information to help you return the server to the same state that the customer typically uses. Examples follow:

- The IPL type that the customer typically uses for the server.
- The IPL mode that is used by the customer on this server.
- How the server is configured or partitioned.

1. Has problem analysis been performed by using the procedures in Beginning problem analysis?
   - **Yes:** Continue with the next step.
   - **No:** Perform problem analysis by using the procedures in Beginning problem analysis.
2. Is the failing server managed by a management console?
   - **Yes:** Continue with step "5" on page 25.
   - **No:** Continue with the next step.
3. Do you have an action plan to perform an isolation procedure?
   - **Yes:** Go to Isolation procedures.
   - **No:** Continue with the next step.
4. Do you have a field replaceable unit (FRU), location code, and an action plan to replace a failing FRU?

- **Yes:** Go to the removal and replacement procedures for the system you are servicing.
- **No:** Go to Part locations and location codes to find the part that you need, and then go to the removal and replacement procedures for the system you are servicing.

This ends the procedure.

5. Is the management console connected and functional?
   - **Yes:** Continue with the next step.
   - **No:** Start the management console and attach it to the server. When the management console is connected and functional, continue with the next step.
6. Were you directed here by support to replace a FRU by using Exchange FRU on the HMC?
   - **Yes:** Go to Exchange FRU.
   - **No:** Continue with the next step.
7. Perform the following steps from the management console that is used to manage the server. During these steps, refer to the service data that was gathered earlier.

   Note: If you are unable to locate the reported problem and there is more than one open problem near the time of the reported failure, use the earliest problem in the list.

   a. In the navigation area, click the **Serviceability** icon 🛠, and then click **Serviceable Events Manager**. The Manage Serviceable Events window is displayed.
   b. From the **Serviceable event status** list, click **Open**.
   c. Select **ALL** for every other selection and click **OK**.
   d. Scroll through the list to determine whether a problem has a status of **Open** and to determine whether it corresponds with the problem reported by the customer.
   e. Do you find the reported problem or an open problem near the time of the reported problem?
      - **Yes:** Continue with the next step.
      - **No:** Go to step "4" on page 24, or if a serviceable event was not found, see the appropriate problem analysis procedure for the operating system you are using.
        - If the server or partition is running the AIX or Linux operating system, see AIX and Linux problem analysis.
        - If the server or partition is running the IBM i operating system, see IBM i problem analysis.
8. To perform a repair operation from the HMC, complete the following steps:

   a. Select the serviceable event that you want to repair, and click **Repair** from the selected menu.
   b. Follow the instructions that are displayed on the HMC.

   After you complete the repair procedure, the system automatically closes the serviceable event. **This ends the procedure.**

Reference information for problem determination

The problem determination reference information is provided as an additional resource for problem detection and analysis when you are directed here by your service representative.

All repair actions should start with "Beginning problem analysis" on page 1 and be followed by "Starting a repair action" on page 24 before you use these tools and techniques for problem determination.

Symptom index

Use this symptom index only when you are guided here by your service representative.

Note: If you were not guided here by your service representative, go to "Beginning problem analysis" on page 1.

Review the symptoms in the left column. Look for the symptom that most closely matches the symptoms on the server that you are troubleshooting. When you find the matching symptom, perform the appropriate action as described in the right column.

Table 5. Determining symptom types

Symptom	What you should do:
You do not have a symptom.	Go to the Starting a repair action.
The symptom or problem is on a server or a partition running IBM i.	Go to "IBM i server or IBM i partition symptoms" on page 26.
The symptom or problem is on a server or a partition running AIX.	Go to "AIX server or AIX partition symptoms" on page 30.
The symptom or problem is on a server or a partition running Linux.	Go to "Linux server or Linux partition symptoms" on page 44.

IBM i server or IBM i partition symptoms

Use the following tables to find the symptom you are experiencing. If you cannot find your symptom, contact your next level of support.

- General symptoms
- Symptoms occurring when the system is not operational
- Symptoms related to a logical partition on a server that has multiple logical partitions
- Obvious physical symptoms
- Time-of-day symptoms

Table 6. General IBM i server or IBM i partition symptoms

Symptom	Service action
You have an intermittent problem or you suspect that the problem is intermittent.	Go to "Intermittent problems" on page 105.
DST/SST functions are available on the logical partition console and: • The customer reports reduced system function. • There is a server performance problem. • There are failing, missing, or inoperable server resources.	On most servers with logical partitions, it is common to have one or more missing or non-reporting system bus resource's under Hardware Service Manager (see Hardware Service Manager for more information).
Operations Console, or the remote control panel is not working properly.	Contact Software Support.
The system has a processor or memory problem.	Use the Service action log to check for a reference code or any failing items. See Service action log for instructions, replacing any hardware FRUs if necessary.

Table 6. General IBM i server or IBM i partition symptoms (continued)

Symptom	**Service action**
The system has detected a bus problem. An SRC of the form B600 69*xx* or B700 69*xx* will be displayed on the control panel or management console.	Go to Service action log.

Table 7. Symptoms occurring when the system is not operational

Symptom	**Service action**
A bouncing or scrolling ball (moving row of dots) remains on the operator panel display, or the operator panel display is filled with dashes or blocks.	Verify that the operator panel connections to the system backplane are connected and properly seated. Also, reseat the Service Processor card. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Operator panel assembly. 2. Service processor.

Table 7. Symptoms occurring when the system is not operational (continued)

Symptom	**Service action**
You have a blank display on the operator panel. Other LEDs on the operator panel appear to behave normally.	Verify that the operator panel connections to the system backplane are connected and properly seated. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Control (operator) panel assembly. 2. Service processor.
There is an IPL problem, the system attention light is on, and blocks of data appear for 5 seconds at a time before moving to the next block of data for 5 seconds, and so on until 5 seconds of a blank control panel is displayed at which time the cycle repeats.	These blocks of data are functions 11 through 20. The first data block after the blank screen is function 11, the second block is function 12, and so on. Use this information to fill out the Problem reporting forms. Then go to Reference codes.
You have a power problem, the system or an attached unit will not power on or will not power off, or there is a 1*xxx*-*xxxx* reference code.	Go to Power problems.
There is an SRC in function 11.	Look up the reference code (see Reference codes).
There is an IPL problem.	Go to “IPL problems” on page 111.
There is a `Device Not Found` message during an installation from an alternate installation device.	Go to TUPIP06.

Table 8. Symptoms related to a logical partition on a server that has multiple logical partitions

Symptom	**Service action**
The console is not working for a logical partition.	See Recovering when the console does not show a sign-on display or a menu with a command line.

Table 8. Symptoms related to a logical partition on a server that has multiple logical partitions (continued)

Symptom	**Service action**
• There is an SRC on the panel of an I/O expansion unit owned by a logical partition. • You suspect a power problem with resources owned by a logical partition. • There is an IPL problem with a logical partition and there is an SRC on the management console. • The logical partition's operations have stopped or the partition is in a loop and there is an SRC on the management console.	• Search Service Focal Point for a serviceable event. • If you do not find a serviceable event in Service Focal Point, then record the partition's SRC from the Operator Panel Values field in the management console. 1. In the navigation area, click the **Resources** icon , and then select **All Systems**. 2. In the content pane, select the required system or click on the server name and then select the required partition. 3. Use that SRC and look up the reference code. For more information, see Reference codes.
The logical partition's console is functioning, but the state of the partition in the management console is "Failed" or "Unit Attn" and there is an SRC.	Use the logical partition's SRC. From the partition's console search for that SRC in the partition's Service Action Log. See Service action log.
There is an IPL problem with a logical partition and there is no SRC displayed in the management console.	Perform the following to look for the panel value for the partition in the management console. 1. In the navigation area, click the **Resources** icon , and then select **All Systems**. 2. In the content pane, select the required system or click on the server name and then select the required partition. 3. In the navigation area, click **Serviceability** > **Reference Code Log** and view the codes. 4. When finished, click **Cancel**. Go to Reference codes. If no reference code can be found, contact your next level of support.
The partition's operations have stopped or the partition is in a loop and there is no SRC displayed on the management console.	Perform function 21 from the management console. If this fails to resolve the problem, contact your next level of support.
One or more of the following was reported: • There is a system reference code or message on the logical partition's console. • The customer reports reduced function in the partition. • There is a logical partition performance problem. • There are failing, missing, or inoperable resources.	From the partition's console search the partition's Service Action Log. Go to Service action log. **Note:** On most systems with logical partitions, it is common to have one or more "Missing or Non-reporting" system bus resource's under Hardware Service Manager. See Hardware Service Manager for details.

Table 8. Symptoms related to a logical partition on a server that has multiple logical partitions (continued)

Symptom	Service action
There is a `Device Not Found` message during an installation from an alternate installation device.	Go to TUPIP06.
There is a problem with a guest partition. **Note:** These are problems reported from the operating system (other than IBM i) running in a guest partition or reported from the hosting partition of a guest partition.	If there are serviceable events in the logical partition or hosting partition, work on these problems first. If there are no SAL entries in the logical partition and no SAL entries in the hosting partition, contact your next level of support.

Table 9. Obvious physical symptoms

Symptom	Service action
A power indicator light, display on the system unit control panel, or an attached I/O unit is not working correctly.	Perform PWR1920.
One or more of the following was reported: • Noise • Smoke • Odor	Go to the system safety inspection procedures for your specific system.
A part is broken or damaged.	Go to the System parts to get the part number. Then go to the remove and replace procedures for your specific system to exchange the part.

Table 10. Time-of-day problems

Symptom	Service action
System clock loses or gains more than 1 second per day when the system is connected to utility power.	Replace the service processor. See symbolic FRU SVCPROC.
System clock loses or gains more than 1 second per day when the system is disconnected from utility power.	Replace the time-of-day battery on the service processor. Go to symbolic FRU TOD_BAT.

AIX server or AIX partition symptoms

Use the following tables to find the symptom you are experiencing. If you cannot find your symptom, contact your next level of support.

Choose the description that best describes your situation:

- You have a service action to complete
- An LED is not operating as expected
- Control (operator) panel problems
- Reference codes
- Management consoles
- There is a display or monitor problem (for example, distortion or blurring)
- Power and cooling problems

- Other symptoms or problems

You have a service action to complete

Symptom	What you should do:
You have an open service event in the service action event log.	Go to Starting a repair action.
You have parts to exchange or a corrective action to complete.	1. Go to the remove and replace procedures for your specific server. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	1. Go to Verifying the repair. 2. Go to the Close of call procedure.
You need to verify correct system operation.	1. Go to Verifying the repair. 2. Go to Close of call procedure.

An LED is not operating as expected

Symptom	What you should do:
The system attention LED on the control panel is on.	Go to Starting a repair action.
The rack indicator LED does not turn on, but a drawer identify LED is on.	1. Make sure that the rack indicator LED is properly mounted to the rack. 2. Make sure that the rack identify LED is properly cabled to the bus bar on the rack and to the drawer identify LED connector. 3. Replace the following parts one at a time: • Rack LED to bus bar cable • LED bus bar to drawer cable • LED bus bar 4. Contact your next level of support.

Control (operator) panel problems

Symptom	What you should do:
01 does not appear in the upper-left corner of the operator panel display after the power is connected and before you press the power-on button. Other symptoms appear in the operator panel display or LEDs before the power on button is pressed.	Go to Power problems.

Symptom	What you should do:
A bouncing or scrolling ball (moving row of dots) remains on the operator panel display, or the operator panel display is filled with dashes or blocks.	Verify that the operator panel connections to the system backplane are connected and properly seated. Also, reseat the service processor card. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Operator panel assembly. 2. Service processor.

Symptom	What you should do:
You have a blank display on the operator panel. Other LEDs on the operator panel appear to behave normally.	Verify that the operator panel connections to the system backplane are connected and properly seated. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Control (operator) panel assembly. 2. Service processor.
You have a blank display on the operator panel. Other LEDs on the operator panel are off.	Go to Power problems.

Reference codes

Symptom	What you should do:
An 8-digit error code is displayed.	Look up the reference code in the Reference codes section of IBM Knowledge Center. **Note:** If the repair for this code does not involve replacing a FRU (for instance, running an AIX command that fixes the problem or changing a hot-pluggable FRU), then update the AIX error log after the problem is resolved by completing the following steps: 1. In the online diagnostics, select **Task Selection** > **Log Repair Action**. 2. Select resource **sysplanar0**. On systems with a fault indicator LED, this changes the fault indicator LED from the fault state to the normal state.

Symptom	What you should do:
The system stops with an 8-digit error code displayed when you boot.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that does **not** begin with 0 or 2.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that begins with 0 or 2.	Record SRN 101-*xxxx*, where *xxxx* is the 4-digit code that is displayed in the control panel. Then, look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.
The system stops and a 3-digit number is displayed on the control panel.	Add 101– to the left of the three digits to create an SRN, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN. If a location code is displayed under the 3-digit error code, look at the location to see whether it matches the failing component that the SRN pointed to. If they do not match, complete the action that is given in the error code table. If the problem still exists, replace the failing component from the location code. If a location code is displayed under the 3-digit error code, record the location code. Record SRN 101-*xxx*, where *xxx* is the 3-digit number that is displayed in the operator panel display, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.

Management consoles

Symptom	What you should do:
The management console cannot be used to manage a managed system, or the connection to the managed system is failing.	If the managed system is operating normally (that is, there are no error codes or other symptoms) the management console might have a problem, or the connection to the managed system might be damaged or incorrectly cabled. Complete the following steps: 1. Check the connections between the management console and the managed system. Correct any cabling errors, if found. If another cable is available, connect it in place of the existing cable and refresh the management console interface. You might have to wait up to 30 seconds for the managed system to reconnect. 2. Verify that any connected management console is connected to the managed system. **Note:** The managed system must have power that is connected, and either be waiting for a power-on instruction (that is, `01` is in the upper-left corner of the operator panel) or be running. If the managed system does not appear in the navigation area of the management console management environment, the management console or the connection to the managed system might be failing. 3. Go to the entry MAP: • Go to: Managing the HMC section. 4. If there is a problem with the service processor card or the system backplane, complete the following steps. • If you cannot fix the problem by using the HMC tests in the Managing the HMC section: a. Replace the service processor card. See the remove and replace procedures for your specific system. b. Replace the system backplane if not already replaced in substep a. See the remove and replace procedures for your specific system.

Symptom	What you should do:
The management console (HMC only) cannot call out by using the attached modem and the customer's telephone line.	If the managed system is operating normally (that is, there are no error codes or other symptoms), the management console might have a problem, or the connection to the modem and telephone line might have a problem. Complete the following steps: 1. Check the connections between the management console, the modem, and the telephone line. Correct any cabling errors, if found. 2. Go to the entry MAP in the Managing your server using the Hardware Management Console section.

There is a display problem (for example, distortion or blurring)

Symptom	What you should do:
All display problems.	1. If you are using the HMC: Go to the Managing the HMC section. 2. If you are using a graphics display, complete the following steps: a. Go to the problem determination procedures for the display. b. If you do not find a problem, complete the following steps, one at a time, until the problem is resolved: 1) Replace the graphics display adapter. See the remove and replace procedures for your specific system. 2) Replace the backplane into which the card is plugged. See the remove and replace procedures for your specific system. 3. If you are using an ASCII terminal, complete the following steps: a. Make sure that the ASCII terminal is connected to S1. b. If problems persist, go to the problem determination procedures for the terminal. c. If you do not find a problem, replace the service processor. See the remove and replace procedures for your specific system.
There appears to be a display problem (distortion, blurring, and so on).	Go to the problem determination procedures for the display.

Power and cooling problems

Symptom	What you should do:
The system does not power on and no error codes are available.	Go to Power problems.
The power LEDs on the operator panel and the power supply do not come on or stay on.	1. Check the service processor error log. 2. Go to Power problems.
The power LEDs on the operator panel and the power supply come on and stay on, but the system does not power on.	1. Check the service processor error log. 2. Go to Power problems.
A rack or a rack-mounted unit will not power on.	1. Check the service processor error log. 2. Go to Power problems.
The cooling fans do not come on, or come on but do not stay on.	1. Check the service processor error log. 2. Go to Power problems.
The system attention LED on the operator panel is on and there is no error code displayed.	1. Check the service processor error log. 2. Go to Power problems.

Other symptoms or problems

Symptom	What you should do:
The system stopped and a code is displayed on the operator panel.	Go to Starting a repair action.
`01` is displayed in the upper-left corner of the operator panel and the fans are off.	The service processor is ready. The system is waiting for power-on. Boot the system. If the boot is unsuccessful, and the system returns to the default display (indicated by `01` in the upper-left corner of the operator panel), go to MAP 0200: Problem determination procedure.
The operator panel displays STBY.	The service processor is ready. The server was shut down by the operating system and is still powered on. This condition can be requested by a privileged system user with no faults. Go to Starting a repair action. **Note:** See the service processor error log for possible operating system fault indications.
All of the system power-on self-test (POST) indicators are displayed on the firmware console, the system pauses and then restarts. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the POST.	Go to Problems with loading and starting the operating system.

Symptom	What you should do:
The system stops and all of the POST indicators are displayed on the firmware console. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the power-on self-test (POST).	Go to Problems with loading and starting the operating system.
The system stops and the message `starting software please wait...`is displayed on the firmware console.	Go to Problems with loading and starting the operating system.
The system does not respond to the password that you entered or the system login prompt is displayed when you boot in service mode.	1. If the password is being entered from HMC: Go to the Managing the HMC. 2. If the password is being entered from a keyboard that is attached to the system, the keyboard or its controller might be faulty. In this case, replace these parts in the following order: a. Keyboard b. Service processor 3. If the password is being entered from an ASCII terminal, use the problem determination procedures for the ASCII terminal. Make sure that the ASCII terminal is connected to S1. If the problem persists, replace the service processor. If the problem is fixed, go to "MAP 0410: Repair checkout" on page 41.
The system stops with a prompt to enter a password.	Enter the password. You cannot continue until a correct password is entered. After you enter a valid password, go to the beginning of this table and wait for one of the other conditions to occur.
The system does not respond when the password is entered.	Go to Step 1020-2.
No codes are displayed on the operator panel within a few seconds of turning on the system. The operator panel is blank before the system is powered on.	Reseat the operator panel cable. If the problem is not resolved, replace in the following order: 1. Operator panel assembly. See the remove and replace procedures for your specific system. 2. Service processor. See the remove and replace procedures for your specific system. If the problem is fixed, go to "MAP 0410: Repair checkout" on page 41. If the problem is still not corrected, go to MAP 0200: Problem determination procedure.

Symptom	What you should do:
The SMS configuration list or boot sequence selection menu shows more SCSI devices that are attached to a controller or an adapter than are actually attached.	A device might be set to use the same SCSI bus ID as the control adapter. Note the ID being used by the controller or adapter (this can be checked or changed through an SMS utility), and verify that no device that is attached to the controller is set to use that ID. If settings do not appear to be in conflict, complete the following steps: 1. Go to MAP 0200: Problem determination procedure. 2. Replace the SCSI cable. 3. Replace the device. 4. Replace the SCSI adapter **Note:** In a **twin-tailed** configuration where there is more than one initiator device (normally another system) attached to the SCSI bus, it might be necessary to use SMS utilities to change the ID of the SCSI controller or adapter.
You have a problem that does not prevent the system from booting. The operator panel is functional and the rack indicator LED operates as expected.	Go to MAP 0200: Problem determination procedure.
All other symptoms.	Go to MAP 0200: Problem determination procedure.
All other problems.	Go to MAP 0200: Problem determination procedure.
You do not have a symptom.	Go to MAP 0200: Problem determination procedure.
You have parts to exchange or a corrective action to complete.	1. Go to Starting a repair action. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	Go to "MAP 0410: Repair checkout" on page 41.
You need to verify correct system operation.	Go to "MAP 0410: Repair checkout" on page 41.

Symptom	What you should do:
The system stopped. A POST indicator is displayed on the system console and an eight-digit error code is not displayed.	If the POST indicator represents: 1. Memory, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 2. Keyboard a. Replace the keyboard. b. Replace the service processor, location: model dependent. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure. 3. Network, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 4. SCSI, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 5. Speaker a. Replace the control panel. The location depends on the model. b. Replace the service processor. The location depends on the model. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The diagnostic operating instructions are displayed.	Go to MAP 0200: Problem determination procedure.
The system login prompt is displayed.	If you are loading the diagnostics from a CD-ROM, you might not have pressed the correct key or you might not have pressed the key soon enough when you were trying to indicate a service mode IPL of the diagnostic programs. If so, try to boot the CD-ROM again and press the correct key. **Note:** Complete the system shutdown procedure before you turn off the system. If you are sure that you pressed the correct key in a timely manner, go to Step 1020-2. If you are loading diagnostics from a Network Installation Management (NIM) server, check for the following items: • The bootlist on the client might be incorrect. • Cstate on the NIM server might be incorrect. • Network problems might be preventing you from connecting to the NIM server. Verify the settings and the status of the network. If you continue to have problems see Problems with loading and starting the operating system and follow the steps for network boot problems.

Symptom	What you should do:
The System Management Services (SMS) menu is displayed when you were trying to boot stand-alone AIX diagnostics.	If you are loading diagnostics from the CD-ROM, you might not have pressed the correct key when you were trying to indicate a service mode IPL of the diagnostic programs. If so, try to boot the CD-ROM again and press the correct key. If you are sure that you pressed the correct key, the device or media you are attempting to boot from might be faulty. 1. Try to boot from an alternate boot device that is connected to the same controller as the original boot device. If the boot succeeds, replace the original boot device (for removable media devices, try the media first). If the boot fails, go to Problems with loading and starting the operating system. 2. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The SMS boot sequence selection menu or remote IPL menu does not show all of the bootable devices in the partition or system.	If an AIX or Linux partition is being booted, verify that the devices that you expect to see in the list are assigned to this partition. If they are not, use the management console to reassign the required resources. If they are assigned to this partition, go to Problems with loading and starting the operating system to resolve the problem.

MAP 0410: Repair checkout

Use this MAP to check out the server after a repair is completed.

Purpose of this MAP

Use this MAP to check out the server after a repair is completed.

Note: Only use standalone diagnostics for repair verification when no other diagnostics are available on the system. Standalone diagnostics do not log repair actions.

If you are servicing an SP system, go to the End-of-call MAP in the *SP System Service Guide*.

If you are servicing a clustered system, go to the End of Call MAP in the *Clustered eServer Installation and Service Guide*.

- **Step 0410-1**

  Did you use an AIX diagnostics service aid hot-swap operation to change the FRU?

  No
  Go to Step 0410-2.

  Yes
  Go to Step 0410-4.

- **Step 0410-2**

  Note: If the system backplane or battery has been replaced and you are loading diagnostics from a server over a network, it may be necessary for the customer to set the network boot information for this system before diagnostics can be loaded. The system time and date information should also be set when the repair is completed.

Do you have any FRUs (for example cards, adapters, cables, or devices) that were removed during problem analysis that you want to put back into the system?

No
Go to Step 0410-3.

Yes
Reinstall all of the FRUs that were removed during problem analysis. Go to Step 0410-3.

- **Step 0410-3**

  Is the system or logical partition that you are performing a repair action on running the AIX operating system?

  No
  Go to Step 0410-5.

  Yes
  Go to Step 0410-4.

- **Step 0410-4**

  Does the system or logical partition you are performing a repair action on have AIX installed?

  Note: Answer **No** to this question if you have just replaced a hard disk in the root volume group.

  No
  Go to Step 0410-5.

  Yes
  Go to Step 0410-6.

- **Step 0410-5**

  Run standalone diagnostics from either a CD ROM or from a NIM server.

  Did you encounter any problems?

  No
  Go to Step 0410-14.

  Yes
  Go to MAP 0020: Problem determination procedure.

- **Step 0410-6**

  1. Power on the system.
  2. Wait until the AIX operating system login prompt displays or until system activity on the operator panel or display apparently has stopped.

  Did the AIX Login Prompt display?

  No
  Go to MAP 0020: Problem determination procedure.

  Yes
  Go to Step 0410-7.

- **Step 0410-7**

  If the **Resource Repair Action** menu is already displayed, go to Step 0410-10; otherwise, complete the following steps:

  1. Log into the operating system either with root authority (if needed, ask the customer to enter the password) or use the CE login.
  2. Enter the `diag -a` command and check for missing resources. Follow any instructions that display. If an SRN displays, suspect a loose card or connection. If no instructions display, no resources were detected as missing. Continue on to Step 0410-8

- **Step 0410-8**

  1. Enter `diag` at the command prompt.

2. Press Enter.
3. Select the **Diagnostics Routines** option.
4. When the DIAGNOSTIC MODE SELECTION menu displays, select **Problem determination**.
5. When the ADVANCED DIAGNOSTIC SELECTION menu displays, select the **All Resources** option or test the FRUs you exchanged, and any devices that are attached to the FRUs you exchanged, by selecting the diagnostics for the individual FRU.

Did the RESOURCE REPAIR ACTION menu (801015) display?

No
Go to Step 0410-9.

Yes
Go to Step 0410-10.

- **Step 0410-9**

  Did the TESTING COMPLETE, no trouble was found menu (801010) display?

  No
  There is still a problem. Go to MAP 0020: Problem determination procedure.

  Yes
  Use the **Log Repair Action** option, if not previously logged, in the TASK SELECTION menu to update the AIX error log. If the repair action was reseating a cable or adapter, select the resource associated with that repair action.

  If the resource associated with your action is not displayed on the resource list, select **sysplanar0**.

  Note: If the check log indicator is on, this will set it back to the normal state.

  Go to Step 0410-12.

- **Step 0410-10**

  When a test is run on a resource in system verification mode, and that resource has an entry in the AIX error log, if the test on the resource was successful, the RESOURCE REPAIR ACTION menu displays.

  After replacing a FRU, you must select the resource for that FRU from the RESOURCE REPAIR ACTION menu. This updates the AIX error log to indicate that a system-detectable FRU has been replaced.

  Note: If the check log indicator is on, this action will set it back to the normal state.

  Complete the following steps:

  1. Select the resource that has been replaced from the RESOURCE REPAIR ACTION menu. If the repair action was reseating a cable or adapter, select the resource associated with that repair action. If the resource associated with your action is not displayed on the resource list, select **sysplanar0**.
  2. Press **Commit** after you make your selections.

  Did another Resource Repair Action (801015) display?

  No
  If the No Trouble Found menu displays, go to Step 0410-12.

  Yes
  Go to Step 0410-11.

- **Step 0410-11**

  The parent or child of the resource you just replaced may also require that you run the RESOURCE REPAIR ACTION service aid on it.

  When a test is run on a resource in system verification mode, and that resource has an entry in the AIX error log, if the test on the resource was successful, the RESOURCE REPAIR ACTION menu displays.

  After replacing that FRU, you must select the resource for that FRU from the RESOURCE REPAIR ACTION menu. This updates the AIX error log to indicate that a system-detectable FRU has been replaced.

Note: If the check log indicator is on, this action will set it back to the normal state.

Complete the following steps:

1. From the RESOURCE REPAIR ACTION menu, select the parent or child of the resource that has been replaced . If the repair action was reseating a cable or adapter, select the resource associated with that repair action. If the resource associated with your action is not displayed on the resource list, select **sysplanar0**.
2. Press COMMIT after you make your selections.
3. If the No Trouble Found menu displays, go to Step 0410-12.

- **Step 0410-12**

  If you changed the service processor or network settings, as instructed in previous MAPs, restore the settings to the value they had prior to servicing the system. If you ran standalone diagnostics from CD-ROM, remove the standalone diagnostics CD-ROM from the system.

  Did you perform service on a RAID subsystem involving changing of the PCI RAID adapter cache card or changing the configuration?

  No
  Go to Step 0410-14.

  Yes
  Go to Step 0410-13.

- **Step 0410-13**

  Use the **Recover Options** selection to resolve the RAID configuration. To do this, complete the following steps:

  1. On the PCI SCSI Disk Array Manager screen, select **Recovery options**.
  2. If a previous configuration exists on the replacement adapter, this must be cleared. Select **Clear PCI SCSI Adapter Configuration**. Press F3.
  3. On the Recovery Options screen, select **Resolve PCI SCSI RAID Adapter Configuration**.
  4. On the Resolve PCI SCSI RAID Adapter Configuration screen, select **Accept Configuration on Drives**.
  5. On the PCI SCSI RAID Adapter selections menu, select the adapter that you changed.
  6. On the next screen, press Enter.
  7. When you get the Are You Sure selection menu, press Enter to continue.
  8. You should get an OK status message when the recover is complete. If you get a `Failed` status message, verify that you selected the correct adapter, then repeat this procedure. When recover is complete, exit the operating system.
  9. Go to Step 0410-14.

- **Step 0410-14**

  If the system you are servicing has a management console, go to the end-of-call procedure for systems with Service Focal Point.

This completes the repair; return the server to the customer.

Linux server or Linux partition symptoms

Use the following tables to find the symptom you are experiencing. If you cannot find your symptom, contact your next level of support.

Choose the description that best describes your situation:

- You have a service action to complete
- An LED is not operating as expected
- Control (operator) panel problems
- Reference codes

- Management consoles
- There is a display or monitor problem (for example, distortion or blurring)
- Power and cooling problems
- Other symptoms or problems

You have a service action to complete

Symptom	What you should do:
You have an open service event in the service action event log.	Go to Starting a repair action.
You have parts to exchange or a corrective action to complete.	1. Go to the remove and replace procedures for your specific system. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	1. Go to Verifying the repair. 2. Go to the Close of call procedure.
You need to verify correct system operation.	1. Go to Verifying the repair. 2. Go to the Close of call procedure.

An LED is not operating as expected

Symptom	What you should do:
The system attention LED on the control panel is on.	Go to “Linux fast-path problem isolation” on page 55.
The rack identify LED does not operate properly.	Go to the “Linux fast-path problem isolation” on page 55.
The rack indicator LED does not turn on, but a drawer identify LED is on.	1. Make sure that the rack indicator LED is properly mounted to the rack. 2. Make sure that the rack identify LED is properly cabled to the bus bar on the rack and to the drawer identify LED connector. 3. Replace the following parts one at a time: • Rack LED to bus bar cable • LED bus bar to drawer cable • LED bus bar 4. Contact your next level of support.

Control (operator) panel problems

Symptom	What you should do:
01 does not appear in the upper-left corner of the operator panel display after the power is connected and before you press the power-on button. Other symptoms appear in the operator panel display or LEDs before the power on button is pressed.	Go to Power problems.

Symptom	What you should do:
A bouncing or scrolling ball (moving row of dots) remains on the operator panel display, or the operator panel display is filled with dashes or blocks.	Verify that the operator panel connections to the system backplane are connected and properly seated. Also, reseat the service processor card. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Operator panel assembly. 2. Service processor.

Symptom	What you should do:
You have a blank display on the operator panel. Other LEDs on the operator panel appear to behave normally.	Verify that the operator panel connections to the system backplane are connected and properly seated. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Control (operator) panel assembly. 2. Service processor.
You have a blank display on the operator panel. Other LEDs on the operator panel are off.	Go to Power problems.

Reference codes

Symptom	What you should do:
An 8-digit error code is displayed.	Look up the reference code in the Reference codes section of IBM Knowledge Center. **Note:** If the repair for this code does not involve replacing a FRU (for instance, running an AIX command that fixes the problem or changing a hot-pluggable FRU), then update the AIX error log after the problem is resolved by completing the following steps: 1. In the online diagnostics, select **Task SelectionLog Repair Action**. 2. Select resource **sysplanar0**. On systems with a fault indicator LED, this changes the "fault indicator" LED from the "fault" state to the "normal" state.

Symptom	What you should do:
The system stops with an 8-digit error code displayed when you boot.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that does **not** begin with 0 or 2.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that begins with 0 or 2 is displayed in the operator panel display.	Record SRN 101-*xxxx*, where *xxxx* is the 4-digit code that is displayed in the control panel, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.
The system stops and a 3-digit number is displayed on the control panel.	Add 101– to the left of the three digits to create an SRN, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN. If a location code is displayed under the 3-digit error code, look at the location to see whether it matches the failing component that the SRN pointed to. If they do not match, complete the action that is given in the error code table. If the problem still exists, then replace the failing component from the location code. If a location code is displayed under the 3-digit error code, record the location code. Record SRN 101-*xxx*, where *xxx* is the 3-digit number that is displayed in the operator panel display, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.

Management consoles

Symptom	**What you should do:**
The management console cannot be used to manage a managed system, or the connection to the managed system is failing.	If the managed system is operating normally (no error codes or other symptoms), the management console might have a problem, or the connection to the managed system might be damaged or incorrectly cabled. Complete the following steps: 1. Check the connections between the management console and the managed system. Correct any cabling errors if found. If another cable is available, connect it in place of the existing cables and refresh the management console interface. You might need to wait up to 30 seconds for the managed system to reconnect. 2. Verify that any connected management console is connected to the managed system by checking the Management Environment of the management console. **Note:** The managed system must have power that is connected and the system running, or waiting for a power-on instruction (`01` is in the upper-left corner of the operator panel). If the managed system does not appear in the Navigation area of the management console Management Environment, the management console or the connection to the managed system might be failing. 3. Go to the Managing the HMC section. 4. There might be a problem with the service processor card or the system backplane. If you cannot fix the problem by using the HMC tests in the Managing the HMC section, complete the following steps, one at a time, until the problem is resolved: a. Replace the service processor card. Refer to the remove and replace procedures for your specific system. b. Replace the management console system backplane. Refer to the remove and replace procedures for your specific system.

Symptom	What you should do:
The management console (HMC only) cannot call out using the attached modem and the customer's telephone line.	If the managed system is operating normally (no error codes or other symptoms), the management console might have a problem, or the connection to the modem and telephone line might have a problem. Complete the following steps: 1. Check the connections between the management console, the modem, and the telephone line. Correct any cabling errors, if found. 2. Go to the Managing your server using the Hardware Management Console section.

There is a display problem (for example, distortion or blurring)

Symptom	What you should do:
All display problems.	1. Go to the Managing the HMC section. 2. If you are using a graphics display, complete the following steps: a. Go to the problem determination procedures for the display. b. If you do not find a problem, complete the following steps, one at a time, until the problem is resolved: 1) Replace the graphics display adapter. Refer to the remove and replace procedures for your specific system. 2) Replace the backplane into which the graphics display adapter is plugged. Refer to the remove and replace procedures for your specific system.
There appears to be a display problem (distortion, blurring, and so on).	Go to the problem determination procedures for the display.

Power and cooling problems

Symptom	What you should do:
The system does not power on and no error codes are available.	Go to Power problems.
The power LEDs on the operator panel and the power supply do not come on or stay on.	1. Check the service processor error log. 2. Go to Power problems.
The power LEDs on the operator panel and the power supply come on and stay on, but the system does not power on.	1. Check the service processor error log. 2. Go to Power problems.
A rack or a rack-mounted unit will not power on.	1. Check the service processor error log. 2. Go to Power problems.

Symptom	What you should do:
The cooling fans do not come on, or come on but do not stay on.	1. Check the service processor error log. 2. Go to Power problems.
The system attention LED on the operator panel is on and no error code is displayed.	1. Check the service processor error log. 2. Go to Power problems.

Other symptoms or problems

Symptom	What you should do:
The system stopped and a code is displayed on the operator panel.	Go to Starting a repair action.
`01` is displayed in the upper-left corner of the operator panel and the fans are off.	The service processor is ready. The system is waiting for power-on. Boot the system. If the boot is unsuccessful, and the system returns to the default display (indicated by `01` in the upper-left corner of the operator panel), go to MAP 0020: Problem determination procedure.
The operator panel displays STBY.	The service processor is ready. The server was shut down by the operating system and is still powered on. This condition can be requested by a privileged system user with no faults. Go to Starting a repair action. **Note:** See the service processor error log for possible operating system fault indications.
All of the system POST indicators are displayed on the firmware console, the system pauses and then restarts. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the power-on self-test (POST).	Go to Problems with loading and starting the operating system.
The system stops and all of the POST indicators are displayed on the firmware console. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the power-on self-test (POST).	Go to Problems with loading and starting the operating system.
The system stops and the message `starting software please wait...`is displayed on the firmware console.	Go to Problems with loading and starting the operating system.

Symptom	What you should do:
The system does not respond to the password that was entered or the system login prompt is displayed when you boot the system in service mode.	1. Go to the Managing the HMC. 2. If the password is being entered from a keyboard that is attached to the system, the keyboard or its controller might be faulty. In this case, replace these parts in the following order: a. Keyboard b. Service processor If the problem is fixed, go to "MAP 0410: Repair checkout" on page 41.
The system stops with a prompt to enter a password.	Enter the password. You cannot continue until a correct password is entered. After you enter a valid password, go to the beginning of this table and wait for one of the other conditions to occur.
The system does not respond when the password is entered.	Go to Step 1020-2.
No codes are displayed on the operator panel within a few seconds of turning on the system. The operator panel is blank before the system is powered on.	Reseat the operator panel cable. If the problem is not resolved, replace in the following order: 1. Operator panel assembly. Refer to the remove and replace procedures for your specific system. 2. Service processor. Refer to the remove and replace procedures for your specific system. If the problem is fixed, go to "MAP 0410: Repair checkout" on page 41. If the problem is still not corrected, go to MAP 0020: Problem determination procedure.
The SMS configuration list or boot sequence selection menu shows more SCSI devices that are attached to a controller/adapter than are actually attached.	A device might be set to use the same SCSI bus ID as the control adapter. Note the ID being used by the controller/adapter (this can be checked or changed through an SMS utility), and verify that no device that is attached to the controller is set to use that ID. If settings do not appear to be in conflict, complete the following steps: 1. Go to MAP 0020: Problem determination procedure. 2. Replace the SCSI cable. 3. Replace the device. 4. Replace the SCSI adapter. **Note:** In a "twin-tailed" configuration where there is more than one initiator device (normally another system) attached to the SCSI bus, it might be necessary to use SMS utilities to change the ID of the SCSI controller or adapter.

Symptom	What you should do:
You suspect a cable problem.	Go to Adapters, Devices and Cables for Multiple Bus Systems.
You have a problem that does not prevent the system from booting. The operator panel is functional and the rack indicator LED operates as expected.	Go to MAP 0020: Problem determination procedure.
All other symptoms.	Go to MAP 0020: Problem determination procedure.
All other problems.	Go to MAP 0020: Problem determination procedure.
You do not have a symptom.	Go to MAP 0020: Problem determination procedure.
You have parts to exchange or a corrective action to complete.	1. Go to Starting a repair action. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	Go to "MAP 0410: Repair checkout" on page 41.
You need to verify correct system operation.	Go to "MAP 0410: Repair checkout" on page 41.
The system stopped. A POST indicator is displayed on the system console and an eight-digit error code is not displayed.	If the POST indicator represents: 1. Memory, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 2. Keyboard a. Replace the keyboard. b. Replace the service processor. The location depends on the model. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure. 3. Network, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 4. SCSI, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 5. Speaker a. Replace the control panel. The location depends on the model. b. Replace the service processor. The location depends on the model. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The diagnostic operating instructions are displayed.	Go to MAP 0020: Problem determination procedure.

Symptom	What you should do:
The system login prompt is displayed.	If you are loading the diagnostics from a CD-ROM, you might not have pressed the correct key or you might not have pressed the key soon enough when you were trying to indicate a service mode IPL of the diagnostic programs. If so, start again at the beginning of this step. **Note:** Complete the system shutdown procedure before you turn off the system. If you are sure that you pressed the correct key in a timely manner, go to Step 1020-2. If you are loading diagnostics from a Network Installation Management (NIM) server, check for the following items: • The bootlist on the client might be incorrect. • Cstate on the NIM server might be incorrect. • There might be network problems that are preventing you from connecting to the NIM server. Verify the settings and the status of the network. If you continue to have problems refer to Problems with loading and starting the operating system and follow the steps for network boot problems.
The System Management Services (SMS) menu is displayed when you were trying to boot stand-alone diagnostics.	If you are loading diagnostics from the CD-ROM, you might not have pressed the correct key when you were trying to indicate a service mode IPL of the diagnostic programs. If so, try to boot the CD-ROM again and press the correct key. If you are sure that you pressed the correct key, the device or media you are attempting to boot from might be faulty. 1. Try to boot from an alternate boot device that is connected to the same controller as the original boot device. If the boot succeeds, replace the original boot device (for removable media devices, try the media first). If the boot fails, go to problems with loading and starting the operating system. 2. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The SMS boot sequence selection menu or remote IPL menu does not show all of the bootable devices in the partition or system.	If an AIX or Linux partition is being booted, verify that the devices that you expect to see in the list are assigned to this partition. If they are not, use the management console to reassign the required resources. If they are assigned to this partition, go to Problems with loading and starting the operating system to resolve the problem.

Linux fast-path problem isolation

Use this information to help you isolate a hardware problem when you use the Linux operating system.

Linux fast path table

Locate the problem in the following table and then go to the action indicated for the problem.

Symptoms	Action
You have an eight-digit reference code.	Go to Reference codes, and do the listed action for the eight-digit reference code.
You are trying to isolate a problem on a Linux server or a partition that is running Linux operating system.	**Note:** This procedure is used to help display an eight-digit reference code by using system log information. Before you use this procedure, if you are having a problem with a media device such as a tape or DVD-ROM drive, continue through this table and follow the actions for the appropriate device. Go to "Linux problem isolation procedure" on page 58.
You suspect a problem with your server but you do not have any specific symptom.	Go to MAP 0020: Problem determination procedure for problem determination procedures.
You need to run the eServer™ stand-alone diagnostics.	Go to AIX and Linux diagnostics and service aids
SRNs	
You have an SRN.	Look up the SRN in the Service request numbers and do the listed action.
An SRN is displayed when you run the eServer stand-alone diagnostics.	1. Record the SRN and location code. 2. Look up the SRN in the Service request numbers and do the listed action.
Tape Drive Problems	
You suspect a tape drive problem.	1. Refer to the tape drive documentation and clean the tape drive. 2. Refer to the tape drive documentation and do any listed problem determination procedures. 3. Go to MAP 0020: Problem determination procedure for problem determination procedures. **Note:** Information on tape cleaning and tape-problem determination is normally either in the tape drive operator guide or the system operator guide.
Optical Drive Problems	
You suspect an optical drive problem.	1. Refer to the optical documentation and do any listed problem determination procedures. 2. Before you service an optical drive, ensure that it is not in use and that the power connector is correctly attached to the drive. If the load or unload operation does not function, replace the optical drive. 3. Go to MAP 0020: Problem determination procedure for problem determination procedures. **Note:** If the optical drive has its own user documentation, follow any problem determination for the optical drive.

Symptoms	Action
SCSI Disk Drive Problems	
You suspect a disk drive problem. Disk problems are logged in the error log and are analyzed when the stand-alone disk diagnostics are run in problem determination mode. Problems are reported if the number of errors is above defined thresholds.	Go to MAP 0020: Problem determination procedure for problem determination procedures.
Token-Ring Problems	
You suspect a token-ring adapter or network problem.	1. Check with the network administrator for known problems. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures.
Ethernet Problems	
You suspect an Ethernet adapter or network problem.	1. Check with the network administrator for known problems. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures.
Display Problems	
You suspect a display problem.	1. If your display is connected to a KVM switch, go to Troubleshooting the keyboard, video, and mouse (KVM) switch for the 1x8 and 2x8 console manager. If you are still having display problems after you complete the KVM switch procedures, come back here and continue with step 2. 2. Go to the Managing your server using the Hardware Management Console section. 3. If you are using a graphics display, complete the following steps: a. Go to the problem determination procedures for the display. b. If you do not find a problem, complete the following steps, one at a time, until the problem is resolved: 1) Replace the graphics display adapter. Refer to the remove and replace procedures for your specific system. 2) Replace the backplane into which the graphics display adapter is plugged. Refer to the remove and replace procedures for your specific system.
Keyboard or Mouse	

Symptoms	Action
You suspect a keyboard or mouse problem.	If your keyboard is connected to a KVM switch, go to Troubleshooting the keyboard, video, and mouse (KVM) switch for the 1x8 and 2x8 console manager. If you are still having keyboard problems after you complete the KVM switch procedures, come back here and continue to the next paragraph. Go to MAP 0020: Problem determination procedure for problem determination procedures. If you are unable to run diagnostics because the system does not respond to the keyboard, replace the keyboard or system backplane. **Note:** If the problem is with the keyboard, it might be caused by the mouse device. To check, unplug the mouse and then recheck the keyboard. If the keyboard works, replace the mouse.
System Messages	
A System Message is displayed.	1. If the message describes the cause of the problem, attempt to correct it. 2. Look for another symptom to use.
System Hangs or Loops when you run the OS or Diagnostics	
The system hangs in the same application.	Suspect the application. To check the system, complete the following steps: 1. Power off the system. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures. 3. If an SRN is displayed at anytime, record the SRN and location code. 4. Look up the SRN in the Service request numbers and do the listed action.
The system hangs in various applications.	1. Power off the system. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures. 3. If an SRN is displayed at anytime, record the SRN and location code. 4. Look up the SRN in the Service request numbers and do the listed action.
The system hangs when you run diagnostics.	Replace the resource that is being tested.
Exchanged FRUs Did Not Fix the Problem	
A FRU or FRUs you exchanged did not fix the problem.	Go to MAP 0020: Problem determination procedure for problem determination procedures.
You Cannot Find the Symptom in This Table	
All other problems.	Go to MAP 0020: Problem determination procedure for problem determination procedures.

Linux problem isolation procedure
Use this procedure when servicing a Linux partition or a server that has Linux as its only operating system.

About this task

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
  - For AC power, disconnect all power cords from their AC power source.
  - For racks with a DC power distribution panel (PDP), disconnect the customer’s DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected.
  - For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
  - For racks with a DC power distribution panel (PDP), connect the customer’s DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements.
- Do not continue with the inspection if any unsafe conditions are present.
- Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.

DANGER:

- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect:

  1. Turn off everything (unless instructed otherwise).
  2. For AC power, remove the power cords from the outlets.
  3. For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source.
  4. Remove the signal cables from the connectors.
  5. Remove all cables from the devices.

To Connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the devices.
3. Attach the signal cables to the connectors.
4. For AC power, attach the power cords to the outlets.
5. For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP.
6. Turn on the devices.

Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

These procedures define the steps to take when servicing a Linux partition or a server that has Linux as its only operating system.

Before continuing with this procedure it is recommended that you review the additional software available to enhance your Linux solutions. See Service and productivity tools for PowerLinux servers (http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags).

Note: If the server is attached to a management console, the various codes that might display on the management console are all listed as reference codes by Service Focal Point (SFP). Use the following table to help you identify the type of error information that might be displayed when you are using this procedure.

Number of digits in reference code	Reference code	Name or code type
Any	Contains # (number sign)	Menu goal
Any	Contains - (hyphen)	Service request number (SRN)
5	Does not contain # or -	SRN
8	Does not contain # or -	system reference code (SRC)

Procedure

1. Is the server managed by a management console that is running Service Focal Point (SFP)?

   No
   Go to step “3” on page 59.

   Yes
   Go to step “2” on page 59.
2. Servers with Service Focal Point

   Look at the service action event log in SFP for errors. Focus on those errors with a timestamp near the time at which the error occurred. Follow the steps indicated in the error log entry to resolve the problem. If the problem is not resolved, continue with step “3” on page 59.
3. Look for and record all reference code information or software messages on the operator panel and in the service processor error log (which is accessible by viewing the ASMI menus).
4. Choose a Linux partition that is running correctly (preferably the partition with the problem).

   Is Linux usable in any partition with Linux installed?

   No
   Go to step “10” on page 60.

   Yes
   Go to step “5” on page 59.
5. Diagnose the RTAS events. For instructions, see Diagnosing RTAS events.

6. Record any RTAS events found in the Linux system log

   If the system is configured with more than one logical partition with Linux installed, repeat step “5” on page 59 and step “6” on page 60 for all logical partitions that have Linux installed.
7. Examine the Linux boot (IPL) log by logging in to the system as the root user and entering the following command:

   ```
   cat /var/log/boot.msg |grep RTAS |more
   ```

   Linux boot (IPL) error messages are logged into the **boot.msg** file under **/var/log**. An example of the Linux boot error log:

   ```
   RTAS daemon started
   RTAS: -------- event-scan begin --------
   RTAS: Location Code: U0.1-F3
   RTAS: WARNING: (FULLY RECOVERED) type: SENSOR
   RTAS: initiator: UNKNOWN target: UNKNOWN
   RTAS: Status: bypassed new
   RTAS: Date/Time: 20020830 14404000
   RTAS: Environment and Power Warning
   RTAS: EPOW Sensor Value: 0x00000001
   RTAS: EPOW caused by fan failure
   RTAS: -------- event-scan end ----------
   ```
8. Record any RTAS events found in the Linux boot (IPL) log in step “7” on page 60.

   Ignore all other events in the Linux boot (IPL) log. If the system is configured with more than one logical partition with Linux installed, repeat step “7” on page 60 and step “8” on page 60 for all logical partitions that have Linux installed.
9. Record any extended data found in the Linux system log in Step “5” on page 59 or the Linux boot (IPL) log in step “7” on page 60.

   Note: The lines in the Linux extended data that begin with `<4>RTAS: Log Debug: 04` contain the reference code listed in the next 8 hexadecimal characters. In the previous example, `4b27 26fb` is a reference code. The reference code is also known as word 11. Each 4 bytes after the reference code in the Linux extended data is another word (for example, `04a0 0011` is word 12, and `702c 0014` is word 13, and so on).

   If the system is configured with more than one logical partition with Linux installed, repeat step “9” on page 60 for all logical partitions that have Linux installed.
10. Were any reference codes or checkpoints recorded in steps “3” on page 59, “6” on page 60, “8” on page 60, or “9” on page 60?

    No

    Go to step “11” on page 60.

    Yes

    Go to the Linux fast-path problem isolation with each reference code that was recorded. Perform the indicated actions one at a time for each reference code until the problem has been corrected. If all recorded reference codes have been processed and the problem has not been corrected, go to step “11” on page 60.
11. If no additional error information is available and the problem has not been corrected, complete the following steps:

    a) Shut down the system.

    b) If a management console is not attached, see Managing your server using the Advanced System Management Interface for instructions to access the ASMI.

    Note: The ASMI functions can also be accessed by using a personal computer connected to system port 1.

    You need a personal computer capable of connecting to system port 1 on the system unit. (The Linux login prompt cannot be seen on a personal computer connected to system port 1.) If the ASMI functions are not otherwise available, use the following procedure:

    1) Attach the personal computer and cable to system port 1 on the system unit.

2) With 01 displayed in the operator panel, press a key on the virtual terminal on the personal computer. The service ASMI menus are available on the attached personal computer.

3) If the service processor menus are not available on the personal computer, perform the following steps:

a) Examine and correct all connections to the service processor.

b) Replace the service processor.

Note: The service processor might be contained on a separate card or board; in some systems, the service processor is built into the system backplane. Contact your next level of support for help before replacing a system backplane.

c) Examine the service processor error log.

Record all reference codes and messages written to the service processor error log. Go to step “12” on page 61.

12. Were any reference codes recorded in step “11” on page 60?

No

Go to step “20” on page 63.

Yes

Go to the Linux fast-path problem isolation with each reference code or symptom you have recorded. Perform the indicated actions, one at a time, until the problem has been corrected. If all recorded reference codes have been processed and the problem has not been corrected, go to “20” on page 63.

13. Reboot the system and bring all partitions to the login prompt.

If Linux is not usable in all partitions, go to step “17” on page 62.

14. Use the `lscfg` command to list all resources assigned to all partitions.

Record the adapter and the partition for each resource.

15. To determine whether any devices or adapters are missing, compare the list of partition assignments, and resources found, to the customer's known configuration. Record the location of any missing devices.

Also record any differences in the descriptions or the locations of devices.

You may also compare this list of resources that were found to an earlier version of the device tree as follows:

Note: At the Linux command prompt, type `vpdupdate`, and press Enter. The device tree is stored in the `/var/lib/lsvpd/` directory in a file with the file name device-tree-YYYY-MM-DD-HH:MM:SS, where YYYY is the year, MM is the month, DD is the day, and HH, MM, and SS are the hour, minute and second, respectively, of the date of creation.

- At the command line, type the following:

```
cd /var/lib/lsvpd/
```

- At the command line, type the following:

```
lscfg -vpz /var/lib/lsvpd/<file_name>
```

Where, *<file_name>* is the .gz file name that contains the database archive.

The **diff** command offers a way to compare the output from a current **lscfg** command to the output from an older **lscfg** command. If the files names for the current and old device trees are **current.out** and **old.out**, respectively, type: `diff old.out current.out`. Any lines that exist in the old, but not in the current will be listed and preceded by a less-than symbol (<). Any lines that exist in the current, but not in the old will be listed and preceded by a greater-than symbol (>). Lines that are the same in both files are not listed; for example, files that are identical will produce no output from the diff command. If the location or description changes, lines preceded by both < and > will be output.

If the system is configured with more than one logical partition with Linux installed, repeat “14” on page 61 and “15” on page 61 for all logical partitions that have Linux installed.

16. Was the location of one and only one device recorded in “15” on page 61?

    No

    If you previously answered Yes to step “16” on page 62, return the system to its original configuration. **This ends the procedure**.

    Go to MAP 0410: Repair checkout.

    If you did not previously answer Yes to step “16” on page 62, go to step “17” on page 62.

    Yes

    Complete the following steps one at a time. Power off the system before each step. After each step, power on the system and go to step “13” on page 61.

    a. Check all connections from the system to the device.
    b. Replace the device (for example, tape or DASD).
    c. If applicable, replace the device backplane.
    d. Replace the device cable.
    e. Replace the adapter.
       - If the adapter resides in an I/O drawer, replace the I/O backplane.
       - If the device adapter resides in the CEC, replace the I/O riser card, or the CEC backplane in which the adapter is plugged.
    f. Call service support. Do not go to step “13” on page 61.

17. Does the system appear to stop or hang before reaching the login prompt or did you record any problems with resources in step “15” on page 61?

    Note: If the system console or VTERM window is always blank, choose NO. If you are sure the console or VTERM is operational and connected correctly, answer the question for this step.

    No

    Go to step “18” on page 62.

    Yes

    There may be a problem with an I/O device. Go to PFW1542: I/O problem isolation procedure. When instructed to boot the system, boot a full system partition.

18. Boot the eServer standalone diagnostics, refer to Running the online and stand-alone diagnostics.

    Run diagnostics in problem determination mode on all resources. Be sure to boot a full system partition. Ensure that diagnostics were run on all known resources. You may need to select each resource individually and run diagnostics on each resource one at a time.

    Did standalone diagnostics find a problem?

    No

    Go to step “22” on page 63.

    Yes

    Go to the Reference codes and perform the actions for each reference code you have recorded. For each reference code not already processed in step “16” on page 62, repeat this action until the problem has been corrected. Perform the indicated actions, one at a time. If all recorded reference codes have been processed and the problem has not been corrected, go to step “22” on page 63.

19. Does the system have Linux installed on one or more partitions?

    No

    Return to the Start a repair action.

    Yes

    Go to step “3” on page 59.

20. Were any location codes recorded in steps “3” on page 59, “6” on page 60, “8” on page 60, “9” on page 60, “10” on page 60, or “11” on page 60?

    No

    Go to step “13” on page 61.

    Yes

    Replace, one at a time, all parts whose location code was recorded in steps “3” on page 59, “6” on page 60, “8” on page 60, “9” on page 60, “10” on page 60, or “11” on page 60 that have not been replaced. Power off the system before replacing a part. After replacing the part, power on the system to check if the problem has been corrected. Go to step “21” on page 63 when the problem has been corrected, or all parts in the location codes list have been replaced.

21. Was the problem corrected in step “20” on page 63?

    No

    Go to step “13” on page 61.

    Yes

    Return the system to its original configuration. **This ends the procedure**.

    Go to MAP 0410: Repair checkout.

22. Were any other symptoms recorded in step “3” on page 59?

    No

    Call support.

    Yes

    Go to the Start a repair action with each symptom you have recorded. Perform the indicated actions for all recorded symptoms, one at a time, until the problem has been corrected. If all recorded symptoms have been processed and the problem has not been corrected, call your next level of support.

Detecting problems

Provides information on using various tools and techniques to detect and identify problems.

IBM i problem determination procedures

There are several tools you can use to determine a problem with an IBM i system or partition.

These include:

Searching the service action log

Use this procedure to search for an entry in the service action log (SAL) that matches the time, reference code, or resource of the reported problem.

Procedure

1. On the command line, enter the Start System Service Tools (STRSST) command. If you cannot get to system service tools (SST), use function 21 to get to dedicated service tools (DST).
2. On the Start Service Tools Sign On display, type in a user ID with QSRV authority and password.
3. Select **Start a Service Tool** > **Hardware Service Manager** > **Work with service action log**.
4. On the Select Timeframe display, change the From: **Date** and **Time** fields to a date and time before the customer reported having the problem.
5. Search for an entry that matches one or more conditions of the problem:
    - Reference code
    - Resource
    - Time
    - Failing item list
6. Perform the following actions:
    - Choose **Display the failing item information** to display the service action log entry.

- Use the **Display details** option to display part location information.

  All new entries in the service action log represent problems that require a service action. It might be necessary to handle any problem in the log even if it does not match the original problem symptom.

  The information that is displayed in the date and time fields are the **Date** and **Time** for the first occurrence of the specific reference code for the resource that is displayed during the time range selected.

7. Did you find an entry in the service action log?

   - **Yes:** Continue with the next step.
   - **No:** Go to "Problems with noncritical resources" on page 105. **This ends the procedure.**

8. Is `See the service information system reference code tables for further problem isolation` shown near the top of the display or are there procedures in the field replaceable unit (FRU) list?

   - **Yes:** Perform the following steps:

     a. Go to the list of reference codes and use the reference code that is indicated in the log to find the correct reference code table and unit reference code.

     b. Perform all actions in the Description/Action column before you replace failing items.

        Note: When you replace failing items, use the part numbers and locations that are found in the service action log entry.

        This ends the procedure.

   - **No:** Display the failing item information for the service action log entry. Items at the top of the failing item list are more likely to fix the problem than items at the bottom of the list.

     Notes:

     a. Some failing items must be replaced in groups until the problem is solved.

     b. Other failing items are flagged as mandatory exchanges and must be replaced before the service action is complete, even if the problem appears to be repaired.

     c. Use the **Part Action Code** field in the Service Action Log display to determine whether failing items are to be replaced in groups or as mandatory exchanges.

     d. Unless the **Part Action Code** of a FRU indicates a group or mandatory exchange, exchange the failing items one at a time until the problem is repaired. Use the help function to determine the meaning of part action codes.

     Continue with the next step.

9. Perform the following steps to help resolve the problem:

   a) To display location information, choose the function key for **Additional details**.

      If location information is available, go to Part locations and location codes for the model you are working on to determine what removal and replacement procedure to perform. To turn on the failing item's identify light, use the indicator-on option.

      Note: In some cases where the failing item does not contain a physical identify light, a higher level identify light is activated (for example, the backplane or the unit that contains the failing item). Use the location information to locate the actual failing item.

   b) If the failing item is Licensed Internal Code, contact your next level of support for the correct fix to apply.

10. After you exchange an item, perform the following steps:

    a) Go to Verifying a repair.

    b) If the failing item indicator was turned on during the removal and replacement procedure, use the indicator-off option to turn off the indicator.

    c) If all problems are resolved for the partition, use the Acknowledge all errors function at the bottom of the service action log display.

   d) Close the log entry by selecting **Close a NEW entry** on the Service Action Log Report display. **This ends the procedure.**

Using the product activity log

This procedure can help you learn how to use the product activity log (PAL).

Procedure

1. To locate a problem, find an entry in the product activity log for the symptom you are seeing.
   a) On the command line, enter the Start System Service Tools (SST) command:

   ```
   STRSST
   ```

   If you cannot get to SST, select DST.

   Note: Do not perform an IPL on the system or partition to get to DST.

   b) On the Start Service Tools Sign On display, type a user ID with service authority and password.

   c) From the System Service Tools display, select **Start a Service Tool** > **Product activity log** > **Analyze log**.

   d) On the Select Subsystem Data display, select the option to view **All Logs**.

   Note: You can change the From: and To: Dates and Times from the 24-hour default if the time that the customer reported having the problem was more than 24 hours ago.

   e) Use the defaults on the Select Analysis Report Options display by pressing the Enter key.

   f) Search the entries on the Log Analysis Report display.

   Note: For example, a 6380 Tape Unit error would be identified as follows:

   System Reference Code: 6380CC5F
   Class: Perm
   Resource Name: TAP01

2. Find an SRC from the product activity log that best matches the time and type of the problem the customer reported.

   Did you find an SRC that matches the time and type of problem the customer reported?

   Yes: Use the SRC information to correct the problem. **This ends the procedure.**
   No: Contact your next level of support. **This ends the procedure**.

Using the problem log

Use this procedure to find and analyze a problem log entry that relates to the problem reported.

About this task

Note: For on-line problem analysis (WRKPRB), ensure that you are logged on with QSRV authority. During problem isolation, this will allow access to test procedures that are not available under any other log-on.

Procedure

1. On the command line, enter the Work with Problems command:

   ```
   WRKPRB
   ```

   Note: Use F4 to change the WRKPRB parameters to select and sort on specific problem log entries that match the problem. Also, F11 displays the dates and times the problems were logged by the system.

   Was an entry that relates to the problem found?

   Note: If the WRKPRB function was not available answer NO.

   Yes: Continue with the next step.

   No: Go to Problems with noncritical resources. **This ends the procedure**.

2. Select the problem entry by moving the cursor to the problem entry option field and entering option 8 to work with the problem.

   Is Analyze Problem (option 1) available on the Work with Problem display?

   No: Perform the following:

   a. Return to the initial problem log display (F12).
   b. Select the problem entry by moving the cursor to the problem entry option field and selecting the option to display details.
   c. Select the function key to display possible causes.

      Note: If this function key is not available, use the customer reported symptom string for customer perceived information about this problem. Then, go to "Using the product activity log" on page 65.
   d. Use the list of possible causes as the FRU list and go to step "5" on page 66.

   Yes: Run Analyze Problem (option 1) from the Work with Problem display.

   Notes:

   a. For SRCs starting with 6112 or 9337, use the SRC and go to the Reference codes topic.
   b. If the message on the display directs you to use SST (System Service Tools), go to COMIP01.

   Was the problem corrected by the analysis procedure?

   No: Continue with the next step.

   Yes: **This ends the procedure**.

3. Did problem analysis send you to another entry point in the service information?

   No: Continue with the next step.

   Yes: Go to the entry point indicated by problem analysis. **This ends the procedure**.

4. Was the problem isolated to a list of failing items?

   Yes: Continue with the next step.

   No: Go to Problems with noncritical resources. **This ends the procedure**.

5. Exchange the failing items one at a time until the problem is repaired.

   Notes:

   a. For Symbolic FRUs, see Symbolic FRUs.
   b. When exchanging FRUs, go to the remove and replace procedures for your specific system.

   Has the problem been resolved?

   No: Contact your next level of support. **This ends the procedure**.

   Yes: **This ends the procedure**.

Problem determination procedure for AIX or Linux servers or partitions

This procedure helps to produce or retrieve a service request number (SRN) if the customer or a previous procedure did not provide one.

If your server is running AIX or Linux, use one of the following procedures to test the server or partition resources to help you determine where a problem might exist.

If you are servicing a server running the AIX operating system, go to MAP 0020: Problem determination procedure.

If you are servicing a server running the Linux operating system, go to the Linux problem isolation procedure.

System unit problem determination

Use this procedure to obtain a reference code if the customer did not provide you with one, or you are unable to load server diagnostics.

If you are able to load the diagnostics, go to Problem determination procedure for AIX or Linux servers or partitions.

The service processor may have recorded one or more symptoms in its error log. Examine this error log before proceeding (For more information, see Managing your server using the Advanced System Management Interface). The server may have been set up by using the management console. Check the Service Action Event (SAE) log in the Service Focal Point. The SAE log may have recorded one or more symptoms in the Service Focal Point. To avoid unnecessary replacement of the same FRU for the same problem, it is necessary to check the SAE log for evidence of prior service activity on the same subsystem.

The service processor may have been set by the user to monitor system operations and to attempt recoveries. You can disable these actions while you diagnose and service the system. If the system maintenance policies were saved by using the save/restore hardware maintenance policies, all the settings of the service processor (except language) were saved and you can use the same service aid to restore the settings at the conclusion of your service action.

If you disable the service processor settings, note their current settings so that you can restore when you are done.

If the system is set to power on using one of the parameters in the following table, disconnect the modem to prevent incoming signals that could cause the system to power on.

Following are the service processor settings. For more information about the service processor settings, see Managing your server using the Advanced System Management Interface.

Table 11. Service processor settings

Setting	Description
Monitoring (also called surveillance)	From the ASMI menu, expand the **System Configuration**, then click on **Monitoring**. Disable both types of surveillance.
Auto power restart (also called unattended start mode)	From the ASMI menu, expand **Power/Restart Control**, then click on **Auto Power Restart**, and set it to disabled.
Wake on LAN	From the ASMI menu, expand **Wake on LAN**, and set it to disabled
Call out	From the ASMI menu, expand the **Service Aids**, then click on **Call-Home/Call-In Setup**. Set the call-home system port and the call-in system port to disabled.

Step 1020-1

Be prepared to record code numbers to help analyze a problem.

Analyze a failure to load the diagnostic programs

Follow these steps to analyze a failure to load the diagnostic programs.

Note: Be prepared to answer questions regarding the control panel and to perform certain actions based on displayed POST indicators. Observer these conditions.

1. Run diagnostics on any partition. Find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, continue to the next step.
2. Run diagnostics on the failing partition. Find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, continue to the next step.
3. Power off the system.

4. Load the standalone diagnostics in service mode to test the full system partition. For more information, see Running the online and stand-alone diagnostics.
5. Wait until the diagnostics are loaded or the system appears to stop. If you receive an error code or if the system stops before diagnostics are loaded, find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, continue to the next step.
6. Run the standalone diagnostics on the entire system. Find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, call service support for assistance.

Symptom	**Action**
One or more logical partitions does not boot.	a. Check service processor error log. If an error is indicated, go to Starting a repair action. b. Check the Serviceable action event log, go to Starting a repair action. c. Go to Problems with loading and starting the operating system.
The rack identify LED does not operate properly.	Go to Starting a repair action.
The system stopped and a system reference code is displayed on the operator panel.	Go to Starting a repair action.
The system stops with a prompt to enter a password.	Enter the password. You cannot continue until a correct password has been entered. When you have entered a valid password, go to the beginning of this table and wait for one of the other conditions to occur.
The diagnostic operating instructions are displayed.	Go to MAP 0020: AIX or Linux problem determination procedure.
The power good LED does not come on or does not stay on, or you have a power problem.	Go to Power problems.
The system login prompt is displayed.	You may not have pressed the correct key or you may not have pressed the key soon enough when you were to trying to indicate a service mode IPL of the diagnostic programs. If this is the case, start again at the beginning of this step. **Note:** Perform the system shutdown procedure before turning off the system. If you are sure you pressed the correct key in a timely manner, go to Step 1020-2.
The system does not respond when the password is entered.	Go to Step 1020-2.

Symptom	Action
The system stopped. A POST indicator is displayed on the system console and an eight-digit error code is not displayed.	If the POST indicator represents: a. Memory, go to PFW 1548: Memory and processor subsystem problem isolation procedure. b. Keyboard 1) Replace the keyboard cable. 2) Replace the keyboard. 3) Replace the service processor. Location is model dependent. 4) Go to PFW1542: I/O problem isolation procedure. c. Network, go to PFW1542: I/O problem isolation procedure. d. SCSI, go to PFW1542: I/O problem isolation procedure. e. Speaker 1) Replace the operator panel. Location is model dependent. 2) Replace the service processor. Location is model dependent. 3) Go to PFW1542: I/O problem isolation procedure.
The System Management Services menu is displayed.	Go to PFW1542: I/O problem isolation procedure.
All other symptoms.	If you were directed here from the Entry MAP, go to PFW1542: I/O problem isolation procedure. Otherwise, find the symptom in the Starting a repair action.

Step 1020-2

Use this procedure to analyze a keyboard problem.

Find the type of keyboard you are using in the following table; then follow the instructions given in the Action column.

Keyboard Type	Action
Type 101 keyboard (U.S.). Identified by the size of the Enter key. The Enter key is in only one horizontal row of keys.	Record error code M0KB D001; then go to Step 1020-3.
Type 102 keyboard (W.T.). Identified by the size of the Enter key. The Enter key extends into two horizontal rows.	Record error code M0KB D002; then go to Step 1020-3.
Type 106 keyboard. (Identified by the Japanese characters.)	Record error code M0KB D003; then go to Step 1020-3.
ASCII terminal keyboard	Go to the documentation for this type of ASCII terminal and continue with problem determination.

Step 1020-3

Perform the following steps:

1. Find the 8-digit error code in Reference codes.

   Note: If you do not locate the 8-digit code, look for it in one of the following places:

   - Any supplemental service manuals for attached devices
   - The diagnostic problem report screen for additional information
   - The Service Hints service aid
   - The CEREADME file
2. Perform the action listed.

Management console machine code problems

The support organization uses the *pesh* command to look at the management console internal machine code to determine how to fix a machine code problem. Only a service representative or support representative can access this feature.

Launching an xterm shell

About this task

You may need to launch an xterm shell to perform directed support from the support center. This may be required if the support center needs to analyze a system dump in order to better understand machine code operations at the time of a failure. To launch an xterm shell, perform the following:

Procedure

1. Open a terminal by right-clicking the background and selecting **Terminals** > **rshterm**.
2. Type the *pesh* command followed by the serial number of the management console and press Enter.
3. You will be prompted for a password, which you must obtain from your next level of support.

Results

Additional information: “Viewing the management console logs” on page 70.

Viewing the management console logs

The console logs display error and information messages that the console has logged while running commands.

About this task

The service representative can use this information to learn more about what caused an error and how to resolve it. The management console classifies log entries as either an informational message or an error message. Log entries are identified with an *I* or *E*, respectively. The management console lists these log entries chronologically, with the most recent shown at the top of the list.

Use the management console Log to view a record of management console system events. System events are activities that indicate when processes begin and end. These events also indicate whether the attempted action was successful.

To view the HMC log, perform the following:

1. Launch an xterm shell (see “Launching an xterm shell” on page 70).
2. When you have entered the password, use the *showLog* command to launch the HMC log window.

The log includes the following information:

- The event's unique ID code
- The date the event occurred
- The time the event occurred
- The log's type
- The name of the attempted action
- The log's reference code
- The status of the log

View a particular event

About this task

To view a particular event, perform the following steps:

Procedure

1. Select an event by clicking once on it.
2. Press Enter to get to a summary of the log you selected. From here, you must select a Block ID to display. The blocks are listed next to the buttons and include the following options:
   - Standard Data Block
   - Secondary Data Block
   - Microcode Reason / ID Error Information
3. Select the data block you want to view.
4. Press Enter. The extended information shown for the data blocks you selected includes the following items:
   - Program name
   - Current process ID
   - Parent process ID
   - Current thread priority
   - Current thread ID
   - Screen group
   - Subscreen group
   - Current foreground screen process group
   - Current background screen process group

Problem determination procedures

Problem determination procedures are provided by power-on self-tests (POSTs), service request numbers, and maintenance analysis procedures (MAPs). Some of these procedures use the service aids that are described in the user or maintenance information for your system SCSI attachment.

Disk drive module power-on self-tests

The disk drive module Power-on Self-Tests (POSTs) start each time that the module is switched on, or when a Send Diagnostic command is received. They check whether the disk drive module is working correctly. The POSTs also help verify a repair after a Field Replaceable Unit (FRU) has been exchanged.

The tests are POST-1 and POST-2.

POST-1 runs immediately after the power-on reset line goes inactive, and before the disk drive module motor starts. POST-1 includes the following tests:

- Microprocessor
- ROM
- Checking circuits

If POST-1 completes successfully, POST-2 is enabled.

If POST-1 fails, the disk drive module is not configured into the system.

POST-2 runs after the disk drive module motor has started. POST-2 includes the following tests:

- Motor control
- Servo control
- Read and write on the diagnostic cylinder (repeated for all heads)
- Error checking and correction (ECC).

If POST-2 completes successfully, the disk drive module is ready for use with the system.

If POST-2 fails, the disk drive module is not configured into the system.

SCSI card power-on self-tests

The SCSI card Power-On Self Tests (POSTs) start each time power is switched on, or when a Reset command is sent from the using system SCSI attachment. They check only the internal components of the SCSI card; they do not check any interfaces to other FRUs.

If the POSTs complete successfully, control passes to the functional microcode of the SCSI card. This microcode checks all the internal interfaces of the I/O enclosure, and reports failures to the host system.

If the POSTs fail, one of the following events occur:

- The SCSI card check LED and the enclosure check LED come on.
- If the SCSI was configured for high availability using a dual initiator card the error will be reported. However, the functional operation of the enclosure is not affected. For example, the customer still has access to all the disk drive modules.

The failure is reported when:

- the failure occurs at system bring-up time, the host system might detect that the enclosure is missing, and reports an error.
- the failure occurs at any time other than system bring-up time, the hourly health check reports the failure.

Analyzing problems

Use these instructions and procedures to help you determine the cause of the problem.

Problems with loading and starting the operating system (AIX and Linux)

If the system is running partitions from partition standby (LPAR), the following procedure addresses the problem in which one partition does not boot AIX or Linux while other partitions boot successfully and run the operating system successfully.

About this task

It is the customer's responsibility to move devices between partitions. If a device must be moved to another partition to run stand-alone diagnostics, contact the customer or system administrator. If the optical drive must be moved to another partition, all SCSI devices that are connected to that SCSI adapter must be moved because moves are done at the slot level, not at the device level.

Depending on the boot device, a checkpoint might be displayed on the operator panel for an extended period while the boot image is retrieved from the device. This is particularly true for tape and network boot attempts. If you are booting from an optical drive or tape drive, watch for activity on the drive's LED indicator. A flashing LED indicates that the loading of either the boot image or additional information that is required by the operating system that is being booted is still in progress. If the checkpoint is displayed for an extended period and the drive LED is not indicating any activity, there might be a problem with loading the boot image from the device.

Notes:

1. For network boot attempts, if the system is not connected to an active network or if the target server is inaccessible (which can also result from incorrect IP parameters), the system still attempts to boot. Because time-out durations are necessarily long to accommodate retries, the system might appear to be hung. Refer to checkpoint CA00 E174.
2. If the partition hangs with a 4-character checkpoint in the display, the partition must be deactivated, then reactivated before you attempt to reboot.
3. If a BA06 000x error code is reported, the partition is already deactivated and in the error state. Reboot by activating the partition. If the reboot is still not successful, go to step “3” on page 73.

This procedure assumes that a diagnostic CD-ROM and an optical drive from which it can be booted are available, or that diagnostics can be run from a NIM (Network Installation Management) server. Booting the diagnostic image from an optical drive or a NIM server is referred to as running stand-alone diagnostics.

Procedure

1. Is a management console attached to the managed system?

   Yes: Continue with the next step.
   No: Go to step "3" on page 73.
2. Look at the service action event error log on the management console.

   Complete the actions necessary to resolve any open entries that affect devices in the boot path of the partition or that indicate problems with I/O cabling. Then, try to reboot the partition. Does the partition reboot successfully?

   Yes: This ends the procedure.
   No: Continue with the next step.
3. Boot to the SMS main menu. Then, choose from the following options:

   - If you are rebooting a partition from partition standby (LPAR), go to the properties of the partition and select **Boot to SMS**, then activate the partition.
   - If you are rebooting from platform standby, access the ASMI. See Setting up and accessing the ASMI. Select **Power/Restart Control**, then **Power On/Off System**. In the AIX/Linux partition mode boot box, select **Boot to SMS menu** > **Save Settings and Power On**.

   At the SMS main menu, select **Select Boot Options** and verify whether the intended boot device is correctly specified in the boot list. Is the intended load device correctly specified in the boot list?

   - **Yes:** Complete the following steps:

     a. Remove all removable media from devices in the boot list from which you do not want to load the operating system.
     b. If you are attempting to load the operating system from a network, go to step "4" on page 73.
     c. If you are attempting to load the operating system from a disk drive or an optical drive, go to step "7" on page 74.
     d. **No:** Go to step "5" on page 73.
4. If you are attempting to load the operating system from the network, complete the following steps:

   - Verify that the IP parameters are correct.
   - Use the SMS ping utility to attempt to ping the target server. If the ping is not successful, have the network administrator verify the server configuration for this client.
   - Check with the network administrator to ensure that the network is up. Also, ask the network administrator to verify the settings on the server from which you are trying to load the operating system.
   - Check the network cabling to the adapter.

   Restart the partition and try loading the operating system. Does the operating system load successfully?

   Yes: This ends the procedure.
   No: Go to step "7" on page 74.
5. Use the SMS menus to add the intended boot device to the boot sequence.

   Can you add the device to the boot sequence?

   Yes: Restart the partition. **This ends the procedure.**
   No: Continue with the next step.

6. Ask the customer or system administrator to verify that the device you are trying to load from is assigned to the correct partition.

   Then, select **List All Devices** and record the list of bootable devices that displays. Is the device from which you want to load the operating system in the list?

   Yes: Go to step "7" on page 74.
   No: Go to step "10" on page 74.

7. Try to load and run stand-alone diagnostics against the devices in the partition, particularly against the boot device from which you want to load the operating system.

   You can run stand-alone diagnostics from an optical drive or a NIM server. To boot stand-alone diagnostics, follow the detailed procedures in Running the online and stand-alone diagnostics.

   Note: When you attempt to load diagnostics on a partition from partition standby, the device from which you are loading stand-alone diagnostics must be made available to the partition that is not able to load the operating system, if it is not already in that partition. Contact the customer or system administrator if a device must be moved between partitions to load stand-alone diagnostics.

   Did stand-alone diagnostics load and start successfully?

   Yes: Go to step "8" on page 74.
   No: Go to step "14" on page 75.

8. Was the intended boot device present in the output of the option **Display Configuration and Resource List** that is run from the Task Selection menu?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "10" on page 74.

9. Did running diagnostics against the intended boot device result in a **No Trouble Found** message?

   Yes: Go to step "12" on page 75.
   No: Go to the list of service request numbers and complete the repair actions for the SRN reported by the diagnostics. After you complete the repair actions, go to step "13" on page 75.

10. Complete the following actions:

    a) Complete the first item in the action list below. In the list of actions below, choose SCSI or IDE based on the type of device from which you are trying to boot the operating system.

    b) Restart the system or partition.

    c) Stop at the SMS menus and select **Select Boot Options**.

    d) Is the device that was not appearing previously in the boot list now present?

       Yes: Go to Verifying a repair. **This ends the procedure.**
       No: Perform the next item in the action list and then return to step "10.b" on page 74. If no more items are in the action list, go to step "11" on page 75.

    Action list:

    Note: See Part locations and location codes for part numbers and links to exchange procedures.

    a) Verify that the SCSI or IDE cables are properly connected. Also, verify that the device configuration and address jumpers are set correctly.

    b) Choose from the following options:

       - **SCSI boot device**: If you are attempting to boot from a SCSI device, remove all hot-swap disk drives (except the intended boot device, if the boot device is a hot-swap drive).If the boot device is present in the boot list after you boot the system to the SMS menus, add the hot-swap disk drives back in one at a time, until you isolate the failing device.
       - **IDE boot device**: If you are attempting to boot from an IDE device, disconnect all other internal SCSI or IDE devices. If the boot device is present in the boot list after you boot the system to the SMS menus, reconnect the internal SCSI or IDE devices one at a time, until you isolate the failing device or cable.

    c) Replace the SCSI or IDE cables.

d) Replace the SCSI backplane (or IDE backplane, if present) to which the boot device is connected.

e) Replace the intended boot device.

f) Replace the system backplane.

11. Choose from the following options:

    - If the intended boot device is not listed, go to "PFW1548: Memory and processor subsystem problem isolation procedure" on page 81. **This ends the procedure.**
    - If an SRN is reported by the diagnostics, go to the list of service request numbers and follow the action that is listed. **This ends the procedure.**

12. Have you disconnected any other devices?

    Yes: Reinstall each device that you disconnected, one at a time. After you reinstall each device, reboot the system. Continue this procedure until you isolate the failing device. Replace the failing device. Then, go to step "13" on page 75.

    No: Perform an operating system-specific recovery process or reinstall the operating system. **This ends the procedure.**

13. Is the problem corrected?

    Yes: Go to Verifying a repair. **This ends the procedure.**

    No: If the replacement of the indicated FRUs did not correct the problem, or if the previous steps did not address your situation, go to "PFW1548: Memory and processor subsystem problem isolation procedure" on page 81. **This ends the procedure.**

14. Is a SCSI boot failure (where you cannot boot from a SCSI-attached device) also occurring?

    - **Yes:** Go to "PFW1548: Memory and processor subsystem problem isolation procedure" on page 81. **This ends the procedure.**
    - **No:** Continue to the next step.

15. Complete the following actions to determine whether another adapter is causing the problem:

    a) Remove all adapters except the one to which the optical drive is attached and the one used for the console.

    b) Reload the stand-alone diagnostics. Can you successfully reload the stand-alone diagnostics?

    - **Yes:** Complete the following steps:

        1) Reinstall the adapters that you removed (and attach devices as applicable) one at a time. After you reinstall each adapter, try the boot operation again until the problem recurs.

        2) Replace the adapter or device that caused the problem.

        3) Go to Verifying a repair. **This ends the procedure.**

    - **No:** Continue with the next step.

16. The graphics adapter (if installed), optical drive, IDE or SCSI cable, or system board is most likely defective.

    Is a PCI graphics adapter installed in the system?

    Yes: Continue with the next step.

    No: Go to step "18" on page 75.

17. Complete the following steps to determine whether the graphics adapter is causing the problem:

    a) Remove the graphics adapter.

    b) Attach a TTY terminal to the system port.

    c) Try to reload stand-alone diagnostics. Do the stand-alone diagnostics load successfully?

    Yes: Replace the graphics adapter. **This ends the procedure.**

    No: Continue with the next step.

18. Replace the following (if not already replaced), one at a time, until the problem is resolved:

    a) Optical drive

b) IDE or SCSI cable that goes to the optical drive.

c) System board that contains the integrated SCSI or IDE adapters.

If this resolves the problem, go to Verifying a repair. If the problem still persists or if the previous descriptions did not address your particular situation, go to “PFW1548: Memory and processor subsystem problem isolation procedure” on page 81.

This ends the procedure.

PFW1540: Problem isolation procedures

The PFW1540 procedures are used to locate problems in the processor subsystem or I/O subsystem.

If a problem is detected, these procedures help you isolate the problem to a failing unit. Find the symptom in the following table; then follow the instructions given in the Action column.

Problem Isolation Procedures

Symptom/Reference Code/Checkpoint	Action
You have or suspect an I/O card or I/O subsystem failure.You received one of the following SRNs or reference codes: 101-000, 101-517, 101-521, 101-538, 101-551 to 101-557, 101-559 to 101-599, 101-662, 101-727, 101-c32, 101-c33, 101-c70	Go to “PFW1542: I/O problem isolation procedure” on page 77.
You have or suspect a memory or processor subsystem problem. You received the following SRN or reference code: 101-185	Go to “PFW1548: Memory and processor subsystem problem isolation procedure” on page 81.
If you were directed to the PFW1540 procedure by an SRN and that SRN is not listed in this table.	Go to “PFW1542: I/O problem isolation procedure” on page 77.

FRU identify LEDs

Your system is configured with an arrangement of LEDs that help identify various components of the system. These include but are not limited to:

- Rack identify beacon LED (optional rack status beacon)
- Processor subsystem drawer identify LED
- I/O drawer identify LED
- FRU identify LED
- Power subsystem FRUs
- Processor subsystem FRUs
- I/O subsystem FRUs
- I/O adapter identify LED
- DASD identify LED

The identify LEDs are arranged hierarchically with the FRU identify LED at the bottom of the hierarchy, followed by the corresponding processor subsystem or I/O drawer identify LED, and the corresponding rack identify LED to locate the failing FRU more easily. Any identify LED in the system may be flashed; see Managing the Advanced system Management Interface (ASMI).

Any identify LED in the system may also be flashed by using the AIX diagnostic programs task "Identify and Attention Indicators". The procedure to use the AIX diagnostics task "Identify and Attention Indicators" is outlined in "diagnostic and service aids" in Running the online and stand-alone diagnostics.

PFW1542: I/O problem isolation procedure

This I/O problem-determination procedure isolates I/O card and I/O subsystem failures. When I/O problem isolation is complete, all cables and cards that are failing will have been replaced or reseated.

For more information about failing part numbers, location codes, or removal and replacement procedures, see Part locations and location codes (http://www.ibm.com/support/knowledgecenter/POWER9/p9ecs/p9ecs_locations.htm). Select your machine type and model number to see applicable procedures for your system.

Notes:

1. To avoid damage to the system or subsystem components, unplug the power cords before removing or installing a part.
2. This procedure assumes either of the following items:
   - An optical drive is installed and connected to the integrated EIDE adapter, and a stand-alone diagnostic CD-ROM is available.
   - Stand-alone diagnostics can be booted from a NIM server.
3. If a power-on password or privileged-access password is set, you are prompted to enter the password before the stand-alone diagnostic CD-ROM can load.
4. The term POST indicators refers to the device mnemonics that appear during the power-on self-test (POST).
5. The service processor might have been set by the user to monitor system operations and to attempt recoveries. You might want to disable these options while you diagnose and service the system. If these settings are disabled, make notes of their current settings so that they can be restored before the system is turned back over to the customer.

   The following settings may be of interest.

   Monitoring
   (also called surveillance) From the ASMI menu, expand the System Configuration menu, then click **Monitoring**. Disable both types of surveillance.

   Auto power restart
   (also called unattended start mode) From the ASMI menu, expand **Power/Restart Control**, then click **Auto Power Restart**, and set it to disabled.

   Wake on LAN
   From the ASMI menu, expand **Wake on LAN**, and set it to disabled.

   Call Out
   From the ASMI menu, expand the Service Aids menu, then click **Call-Home/Call-In Setup**. Set the call-home system port and the call-in system port to disabled.
6. Verify that the system has not been set to boot to the SMS menus or to the open firmware prompt. From the ASMI menu, expand **Power/Restart Control** to view the menu, then click **Power On/Off System**. The AIX/Linux partition mode boot indicates **Continue to Operating System**.

Use this procedure to locate defective FRUs not found by normal diagnostics. For this procedure, diagnostics are run on a minimally configured system. If a failure is detected on the minimally configured system, the remaining FRUs are exchanged one at a time until the failing FRU is identified. If a failure is not detected, FRUs are added back until the failure occurs. The failure is then isolated to the failing FRU.

Perform the following procedure:

- **PFW1542-1**
  1. Ensure that the diagnostics and the operating system are shut down.
  2. Turn off the power.
  3. Turn on the power.
  4. Insert the stand-alone diagnostic CD-ROM into the optical drive.

  Does the optical drive appear to operate correctly?

No
Go to “Problems with loading and starting the operating system (AIX and Linux)” on page 72.

Yes
Continue to PFW1542-2.

- **PFW1542-2**

 1. When the keyboard indicator is displayed (the word "keyboard"), if the system or partition gets that far in the IPL process, press the 5 key on the firmware console.
 2. If you are prompted to do so, enter the appropriate password.

 Is the "Please define the System Console" screen displayed?

 No
 Continue to PFW1542-3.

 Yes
 Go to PFW1542-4.

- **PFW1542-3**

 The system is unable to boot stand-alone diagnostics.

 Did powering on the system generate a different error code or partition firmware hang from the one that originally sent you to PFW1542?

 No
 If you were sent here by an error code, and the error code did not change as the result of powering on the system, you have a processor subsystem problem. Go to “PFW1548: Memory and processor subsystem problem isolation procedure” on page 81. If you were sent here because the system is hanging on a partition firmware checkpoint, and the hang condition did not change as a result of powering on the system, go to PFW1542-5.

 Yes
 Look up the new error code in the reference code index and perform the listed actions.

- **PFW1542-4**

 The system stopped with the **Please define the System Console** prompt on the system console. Stand-alone diagnostics can be booted. Perform the following steps:

 1. Follow the instructions on the screen to select the system console.
 2. When the DIAGNOSTIC OPERATING INSTRUCTIONS screen is displayed, press Enter.
 3. If the terminal type has not been defined, you must use the option **Initialize Terminal** on the FUNCTION SELECTION menu to initialize the AIX operating system environment before you can continue with the diagnostics. This is a separate operation from selecting the firmware console.
 4. Select **Advanced Diagnostic Routines**.
 5. When the DIAGNOSTIC MODE SELECTION menu displays, select **System Verification** to run diagnostics on all resources.

    Did running diagnostics produce a different symptom?

    No
    Continue with the following sub-step.

    Yes
    Return to the Problem Analysis procedures with the new symptom.
 6. Record any devices missing from the list of all adapters and devices. Continue with this procedure. When you have fixed the problem, use this record to verify that all devices appear when you run system verification.

    Are there any devices missing from the list of all adapters and devices?

No
Reinstall all remaining adapters, if any, and reconnect all devices. Return the system to its original configuration. Go to Verifying a repair.

Yes
The boot attempts that follow will attempt to isolate any remaining I/O subsystem problems with missing devices. Ignore any codes that may display on the operator panel unless stated otherwise. Continue to PFW1542-5.

- **PFW1542-5**

  Are there any adapters in the PCI slots in the base system?

  No
  Go to PFW1542-6.

  Yes
  Go to PFW1542-8.

- **PFW1542-6**

  Replace the system backplane, U*n*-P1. Continue to PFW1542-7.

- **PFW1542-7**

  1. Boot stand-alone diagnostics from CD.
  2. If the "Please define the System Console" screen is displayed, follow directions to select the system console.
  3. Use the Display Configuration and Resource List to list all adapters and attached devices.
  4. Check that all adapters and attached devices are listed.

  Did the "Please define the System Console" screen display and are all attached devices and adapters listed?

  No
  Go to PFW1542-11.

  Yes
  Go to PFW1542-12.

- **PFW1542-8**

  1. If it is not already off, turn off the power.
  2. Label and record the location of any cables attached to the adapters.
  3. Record the slot number of the adapters.
  4. Remove all adapters from slots 1, 2, 3, 4, 5, and 6 in the base system that are not attached to the boot device.
  5. Turn on the power to boot stand-alone diagnostics from CD-ROM.
  6. If the ASCII terminal displays **Enter 0 to select this console**, press the 0 key on the ASCII terminal's keyboard.
  7. If the "Please define the System Console" screen is displayed, follow directions to select the system console.
  8. Use the option **Display Configuration and Resource List** to list all adapters and attached devices.
  9. Check that all adapters and attached devices are listed.

  Did the "Please define the System Console" screen display and are all attached devices and adapters listed?

  No
  Go to PFW1542-11.

  Yes
  Continue to PFW1542-9.

- **PFW1542-9**

  If the "Please define the System Console" screen does display and all adapters and attached devices are listed, the problem is with one of the adapters or devices that was removed or disconnected from the base system.

  1. Turn off the power.
  2. Reinstall one adapter and device that was removed. Use the original adapters in their original slots when reinstalling adapters.
  3. Turn on the power to boot stand-alone diagnostics from the optical drive.
  4. If the "Please define the System Console" screen is displayed, follow the directions to select the system console.
  5. Use the Display Configuration and Resource List to list all adapters and attached devices.
  6. Check that all adapters and attached devices are listed.

  Did the "Please define the System Console" screen display and are all attached devices and adapters listed?

  No
  Continue to PFW1542-10.

  Yes
  Return to the beginning of this step to continue reinstalling adapters and devices.

- **PFW1542-10**

  Replace the adapter you just installed with a new adapter and try the boot to stand-alone diagnostics from CD-ROM again.

  1. If the "Please Define the System Console" screen is displayed, follow directions to select the system console.
  2. Use the option **Display Configuration and Resource List** to list all adapters and attached devices.
  3. Check that all adapters and attached devices are listed.

  Did the "Please define the System Console" screen display and are all attached devices and adapters listed?

  No
  Go to PFW1542-6.

  Yes
  The adapter you just replaced was defective. Go to PFW1542-12.

- **PFW1542-11**

  1. Turn off the power.
  2. Disconnect the base system power cables.
  3. Replace the following parts, one at a time, in the sequence listed:
     a. Optical drive
     b. Removable media backplane and cage assembly
     c. Disk drive backplane and cage assembly
     d. System backplane, location U*n*-P1
     e. Service processor
  4. Reconnect the base system power cables.
  5. Turn on the power.
  6. Boot stand-alone diagnostics from CD.
  7. If the "Please define the System Console" screen is displayed, follow directions to select the system console.

8. Use the option **Display Configuration and Resource List** to list all adapters and attached devices.
9. Check that all adapters and attached devices are listed.

Did the "Please define the System Console" screen display and are all adapters and attached devices listed?

No

Replace the next part in the list and return to the beginning of this step. Repeat this process until a part causes the Please define the System Console screen to be displayed and all adapters and attached devices to be listed. If you have replaced all the items listed above and the Please define the System Console screen does not display or all adapters and attached devices are not listed, check all external devices and cabling. If you do not find a problem, contact your next level of support for assistance.

Yes

Go to PFW1542-12.

- **PFW1542-12**

  The item you just replaced fixed the problem.

  1. Turn off the power.
  2. If a display adapter with keyboard and mouse were installed, reinstall the display adapter, keyboard, and mouse.
  3. Reconnect the tape drive (if previously installed) to the internal SCSI bus cable.
  4. Plug in all adapters that were previously removed but not reinstalled.
  5. Reconnect the I/O subsystem power cables that were previously disconnected.

  Return the system to its original condition. Go to Verifying a repair.

PFW1548: Memory and processor subsystem problem isolation procedure

Use this problem isolation procedure to aid in solving memory and processor problems that are not found by normal diagnostics.

Notes:

1. To avoid damage to the system or subsystem components, unplug the power cords before removing or installing any part.
2. This procedure assumes that either:
   - An optical drive is installed and connected to the integrated EIDE adapter, and a stand-alone diagnostic CD-ROM is available.

     OR
   - Stand-alone diagnostics can be booted from a NIM server.
3. If a power-on password or privileged-access password is set, you are prompted to enter the password before the stand-alone diagnostic CD-ROM can load.
4. The term POST indicators refers to the device mnemonics that appear during the power-on self-test (POST).
5. The service processor might have been set by the user to monitor system operations and to attempt recoveries. You might want to disable these options while you diagnose and service the system. If these settings are disabled, make notes of their current settings so that they can be restored before the system is turned back over to the customer. The following settings may be of interest.

   Monitoring

   (also called surveillance) From the ASMI menu, expand the **System Configuration** menu, then click **Monitoring**. Disable both types of surveillance.

   Auto power restart

   (also called unattended start mode) From the ASMI menu, expand **Power/Restart Control**, then click **Auto Power Restart**, and set it to disabled.

Wake on LAN
From the ASMI menu, expand **Wake on LAN**, and set it to disabled.

Call Out
From the ASMI menu, expand the **Service Aids** menu, then click **Call-Home/Call-In Setup**. Set the call-home system port and the call-in system port to disabled.

6. Verify that the system has not been set to boot to the System Management Services (SMS) menus or to the open firmware prompt. From the ASMI menu, expand **Power/Restart Control** to view the menu, then click **Power On/Off System**. The AIX/Linux partition mode boot should say "Continue to Operating System".
7. The service processor might have recorded one or more symptoms in its error/event log. Use the Advanced System Management Interface (ASMI) menus to view the error/event log.

   - Look for a possible new error that occurred during power on of the system. If there is a new error, and its actions call for a FRU replacement, perform those actions. If this does not resolve the problem, go to PFW1548-1.
   - If powering on the system did not yield a new error code, look at the error that occurred just before the original error. Perform the actions associated with that error. If this does not resolve the problem, go to PFW1548-1.
   - If powering on the system results in the same error code, and there are no error codes before the original error code, go to PFW1548-1.

Perform the following procedure:

- **PFW1548-1**

  1. Ensure that the diagnostics and the operating system are shut down.

     Is the system at "service processor standby", indicated by 01 in the control panel?

     No
     Replace the system backplane. Return to the beginning of this step.

     Yes
     Continue with substep 2.

  2. Turn on the power using either the white button or the ASMI menus.

     If an HMC is attached, does the system reach hypervisor standby as indicated on the management console? If a management console is not attached, does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the Please define the System Console screen displayed?

     No
     Go to PFW1548-3.

     Yes
     Go to PFW1548-2.

  3. Insert the stand-alone diagnostic CD-ROM into the optical drive.

     Note: If you cannot insert the diagnostic CD-ROM, go to PFW1548-2.

  4. When the word *keyboard* is displayed on an ASCII terminal, a directly attached keyboard, or management console, press the number 5 key.
  5. If you are prompted to do so, enter the appropriate password.

     Is the "Please define the System Console" screen displayed?

     No
     Go to PFW1548-2.

     Yes
     Go to PFW1548-14.

- **PFW1548-2**

Insert the stand-alone diagnostic CD-ROM into the optical drive.

Note: If you cannot insert the stand-alone diagnostic CD-ROM, go to step PFW1548-3.

Turn on the power using either the white button or the ASMI menus. (If the stand-alone diagnostic CD-ROM is not in the optical drive, insert it now.) If a management console is attached, after the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the stand-alone diagnostic CD-ROM.

If you are prompted to do so, enter the appropriate password.

Is the "Please define the System Console" screen displayed?

No
Go to PFW1548-3.

Yes
Go to PFW1548-14.

- **PFW1548-3**

1. Turn off the power.
2. If you have not already done so, configure the service processor (using the ASMI menus) with the instructions in note 6 at the beginning of this procedure, then return here and continue.
3. Exit the service processor (ASMI) menus and remove the power cords.
4. Disconnect all external cables (parallel, system port 1, system port 2, keyboard, mouse, USB devices, SPCN, Ethernet, and so on). Also disconnect all of the external cables attached to the service processor except the Ethernet cable going to the management console, if a management console is attached.

Go to the next step.

- **PFW1548-4**

Perform the following steps:

1. Place the drawer into the service position and remove the service access cover.
2. Record the slot numbers of the PCI adapters and I/O expansion cards if present. Label and record the locations of all cables attached to the adapters. Disconnect all cables attached to the adapters and remove all of the adapters.
3. Slide the media or disk drive enclosure out approximately three centimeters.
4. Remove and label the disk drives from the media or disk drive enclosure assembly.
5. Remove all memory DIMMs except for one pair.
6. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
7. Turn on the power using either the management console or the white button.

- **PFW1548-5**

Were any memory DIMMs removed from system backplane?

No
Go to PFW1548-8.

Yes
Go to the next step.

- **PFW1548-6**

1. Turn off the power, and remove the power cords.
2. Replug the memory DIMMs that were removed in PFW1548-4 in their original locations.
3. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
4. Turn on the power using either the management console or the white button.

If a management console is attached, does the managed system reach power on at hypervisor standby as indicated on the management console? If a management console is not attached, does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the Please define the System Console screen displayed?

No:

A memory DIMM in the pair you just replaced in the system is defective. Turn off the power, remove the power cords, and exchange the memory DIMM pair with new or previously removed memory DIMM pair. Repeat this step until the defective memory DIMM pair is identified, or all memory DIMM pairs have been replaced.

If your symptom did not change and all the memory DIMM pairs have been exchanged, call your service support person for assistance. If the symptom changed, check for loose cards and obvious problems.

If you do not find a problem, go to Problem Analysis and follow the instructions for the new symptom.

Yes:

Go to PFW1548-7.1.

- **PFW1548-7.1**

  No failure was detected with this configuration.

  1. Turn off the power and remove the power cords.
  2. Reinstall the next DIMM pair.
  3. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
  4. Turn on the power using either the management console or the white button.

  If a management console is attached, does the managed system reach power on at hypervisor standby as indicated on the management console? If a management console is not attached, does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the Please define the System Console screen displayed?

  No:

  One of the FRUs remaining in the system is defective. Exchange the FRUs (that have not already been changed) in the following order:

  a. Memory DIMMs (if present). Exchange the DIMM pairs, one at a time, with new or previously removed DIMM pairs.
  b. System backplane
  c. Power supplies
  d. Processor modules

  Repeat the FRU replacement steps until the defective FRU is identified or all the FRUs have been exchanged.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis and follow the instructions for the new symptom.

  Yes:

  If all of the processor cards have been reinstalled, go to step PFW1548-8. Otherwise, repeat this step.

- **PFW1548-8**

  1. Turn off the power.
  2. Reconnect the system console.

  Notes:

   a. If an ASCII terminal has been defined as the firmware console, attach the ASCII terminal cable to the S1 connector on the rear of the system unit.
   b. If a display attached to a display adapter has been defined as the firmware console, install the display adapter and connect the display to the adapter. Plug the keyboard and mouse into the keyboard connector on the rear of the system unit.
3. Turn on the power using either the management console or the white button. (If the stand-alone diagnostic CD-ROM is not in the optical drive, insert it now.) If a management console is attached, after the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the stand-alone diagnostic CD-ROM.
4. If the ASCII terminal or graphics display (including display adapter) is connected differently from the way it was previously, the console selection screen appears. Select a firmware console.
5. Immediately after the word *keyboard* is displayed, press the number 1 key on the directly attached keyboard, an ASCII terminal or management console. This activates the system management services (SMS).
6. Enter the appropriate password if you are prompted to do so.

Is the SMS screen displayed?

No

One of the FRUs remaining in the system unit is defective.

If you are using an ASCII terminal, go to the problem determination procedures for the display. If you do not find a problem, replace the system backplane.

Yes

Go to the next step.

- **PFW1548-9**

1. Make sure the stand-alone diagnostic CD-ROM is inserted into the optical drive.
2. Turn off the power and remove the power cords.
3. Use the cam levers to reconnect the disk drive enclosure assembly to the I/O backplane.
4. Reconnect the removable media or disk drive enclosure assembly.
5. Plug in the power cords and wait for 01 in the upper-left corner of the operator panel display.
6. Turn on the power using either the management console or the white button. (If the stand-alone diagnostic CD-ROM is not in the optical drive, insert it now.) If a management console is attached, after the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the stand-alone diagnostic CD-ROM.
7. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or an ASCII terminal keyboard.
8. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No:

One of the FRUs remaining in the system unit is defective.

Exchange the FRUs in the order listed that have not been exchanged.

1. Optical drive
2. Removable media enclosure
3. System backplane

Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis and follow the instructions for the new symptom.

Yes:
Go to the next step.

- **PFW1548-10**

  The system is working correctly with this configuration. One of the disk drives that you removed from the disk drive backplanes may be defective.

  1. Make sure the stand-alone diagnostic CD-ROM is inserted into the optical drive.
  2. Turn off the power and remove the power cords.
  3. Install a disk drive in the media or disk drive enclosure assembly.
  4. Plug in the power cords and wait for the OK prompt to display on the operator panel display.
  5. Turn on the power.
  6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or an ASCII terminal keyboard.
  7. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No
  Exchange the FRUs in the order listed that have not been exchanged.

  1. Last disk drive installed
  2. Disk drive backplane

  Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis and follow the instructions for the new symptom.

  Yes
  Repeat this step with all disk drives that were installed in the disk drive backplane.

  After all of the disk drives have been reinstalled, go to the next step.

- **PFW1548-11**

  The system is working correctly with this configuration. One of the devices that was disconnected from the system backplane may be defective.

  1. Turn off the power and remove the power cords.
  2. Attach a system backplane device (for example: system port 1, system port 2, USB, keyboard, mouse, Ethernet) that had been removed.

     After all of the I/O backplane device cables have been reattached, reattached the cables to the service processor one at a time.
  3. Plug in the power cords and wait for 01 in the upper-left corner on the operator panel display.
  4. Turn on the power using either the management console or the white button. (If the stand-alone diagnostic CD-ROM is not in the optical drive, insert it now.) If a management console is attached, after the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the stand-alone diagnostic CD-ROM.
  5. If the Console Selection screen is displayed, choose the system console.
  6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
  7. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No

The last device or cable that you attached is defective.

To test each FRU, exchange the FRUs in the order listed.

1. Device and cable (last one attached).
2. System backplane

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis and follow the instructions for the new symptom.

Yes

Repeat this step until all of the devices are attached. Go to the next step.

- **PFW1548-12**

The system is working correctly with this configuration. One of the FRUs (adapters) that you removed may be defective.

1. Turn off the power and remove the power cords.
2. Install a FRU (adapter) and connect any cables and devices that were attached to the FRU.
3. Plug in the power cords and wait for the OK prompt to display on the operator panel display.
4. Turn on the power using either the management console or the white button. (If the stand-alone diagnostic CD-ROM is not in the optical drive, insert it now.) If a management console is attached, after the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the stand-alone diagnostic CD-ROM.
5. If the Console Selection screen is displayed, choose the system console.
6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
7. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No

Go to the next step.

Yes

Repeat this step until all of the FRUs (adapters) are installed. Go to Verifying a repair.

- **PFW1548-13**

The last FRU installed or one of its attached devices is probably defective.

1. Make sure the stand-alone diagnostic CD-ROM is inserted into the optical drive.
2. Turn off the power and remove the power cords.
3. Starting with the last installed adapter, disconnect one attached device and cable.
4. Plug in the power cords and wait for the 01 in the upper-left corner on the operator panel display.
5. Turn on the power using either the management console or the white button. (If the stand-alone diagnostic CD-ROM is not in the optical drive, insert it now.) If a management console is attached, after the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the Advanced activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the stand-alone diagnostic CD-ROM.
6. If the Console Selection screen is displayed, choose the system console.
7. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.

8. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No

Repeat this step until the defective device or cable is identified or all devices and cables have been disconnected.

If all the devices and cables have been removed, then one of the FRUs remaining in the system unit is defective.

To test each FRU, exchange the FRUs in the order listed.

1. Adapter (last one installed)
2. System backplane

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis and follow the instructions for the new symptom.

Yes

The last device or cable that you disconnected is defective. Exchange the defective device or cable then go to the next step.

- **PFW1548-14**

1. Follow the instructions on the screen to select the system console.
2. When the DIAGNOSTIC OPERATING INSTRUCTIONS screen is displayed, press Enter.
3. Select **Advanced Diagnostics Routines**.
4. If the terminal type has not been defined, you must use the option **Initialize Terminal** on the FUNCTION SELECTION menu to initialize the diagnostic environment before you can continue with the diagnostics. This is a separate operation from selecting the console display.
5. If the NEW RESOURCE screen is displayed, select an option from the bottom of the screen.

   Note: Adapters and devices that require supplemental media are not shown in the new resource list. If the system has adapters or devices that require supplemental media, select option 1.
6. When the DIAGNOSTIC MODE SELECTION screen is displayed, press Enter.
7. Select **All Resources**. (If you were sent here from step PFW1548-18, select the adapter or device that was loaded from the supplemental media).

Did you get an SRN?

No

Go to step PFW1548-16.

Yes

Go to the next step.

- **PFW1548-15**

Look at the FRU part numbers associated with the SRN.

Have you exchanged all the FRUs that correspond to the failing function codes (FFCs)?

No

Exchange the FRU with the highest failure percentage that has not been changed.

Repeat this step until all the FRUs associated with the SRN have been exchanged or diagnostics run with no trouble found. Run diagnostics after each FRU is exchanged. Go to Verifying a repair.

Yes

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

- **PFW1548-16**

  Does the system have adapters or devices that require supplemental media?

  No
  Go to step the next step.

  Yes
  Go to step PFW1548-18.

- **PFW1548-17**

  Consult the PCI adapter configuration documentation for your operating system to verify that all adapters are configured correctly.

  Go to Verifying a repair.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

- **PFW1548-18**

  1. Select **Task Selection**.
  2. Select **Process Supplemental Media** and follow the on-screen instructions to process the media. Supplemental media must be loaded and processed one at a time.

  Did the system return to the TASKS SELECTION SCREEN after the supplemental media was processed?

  No
  Go to the next step.

  Yes
  Press F3 to return to the FUNCTION SELECTION screen. Go to step PFW1548-14 substep 4.

- **PFW1548-19**

  The adapter or device is probably defective.

  If the supplemental media is for an adapter, replace the FRUs in the following order:

  1. Adapter
  2. System backplane

  If the supplemental media is for a device, replace the FRUs in the following order:

  1. Device and any associated cables
  2. The adapter to which the device is attached

  Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis and follow the instructions for the new symptom.

  Go to Verifying a repair.

  This ends the procedure.

PFW1548: Memory and processor subsystem problem isolation procedure when a management console is attached

This procedure is used to locate defective FRUs not found by normal diagnostics. For this procedure, diagnostics are run on a minimally configured system. If a failure is detected on the minimally configured system, the remaining FRUs are exchanged one at a time until the failing FRU is identified. If a failure is not detected, FRUs are added back until the failure occurs. The failure is then isolated to the failing FRU.

Perform the following procedure:

- **PFW1548-1**

  1. Ensure that the diagnostics and the operating system are shut down.

     Is the system at "service processor standby", indicated by 01 in the control panel?

     No
     Replace the system backplane, location: U*n*-P1. Return to step PFW1548-1.

     Yes
     Continue with substep “2” on page 90.
  2. Turn on the power using either the white button or the ASMI menus.

     Does the system reach hypervisor standby as indicated on the management console?

     No
     Go to PFW1548-3.

     Yes
     Go to PFW1548-2.
  3. Insert the stand-alone diagnostic CD-ROM into the optical drive.

     Note: If you cannot insert the diagnostic CD-ROM, go to PFW1548-2.
  4. When the word *keyboard* is displayed on an ASCII terminal, a directly attached keyboard, or management console, press the number 5 key.
  5. If you are prompted to do so, enter the appropriate password.

     Is the "Please define the System Console" screen displayed?

     No
     Go to PFW1548-2.

     Yes
     Go to PFW1548-14.
- **PFW1548-2**

  Insert the stand-alone diagnostic CD-ROM into the optical drive.

  Note: If you cannot insert the diagnostic CD-ROM, go to step PFW1548-3.

  Turn on the power using either the white button or the ASMI menus. (If the diagnostic CD-ROM is not in the optical drive, insert it now.) After the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the diagnostic CD-ROM.

  If you are prompted to do so, enter the appropriate password.

  Is the "Please define the System Console" screen displayed?

  No
  Go to PFW1548-3.

  Yes
  Go to PFW1548-14.
- **PFW1548-3**

  1. Turn off the power.
  2. If you have not already done so, configure the service processor (using the ASMI menus), follow the instructions in note 6 located in “PFW1548: Memory and processor subsystem problem isolation procedure” on page 81 and then return here and continue.
  3. Exit the service processor (ASMI) menus and remove the power cords.
  4. Disconnect all external cables (parallel, system port 1, system port 2, keyboard, mouse, USB devices, SPCN, Ethernet, and so on). Also disconnect all of the external cables attached to the service processor except the Ethernet cable going to the management console.

Go to the next step.

- **PFW1548-4**

  1. If this is a deskside system, remove the service access cover. If this is a rack-mounted system, place the drawer into the service position and remove the service access cover. Also remove the front cover.
  2. Record the slot numbers of the PCI adapters and I/O expansion cards if present. Label and record the locations of all cables attached to the adapters. Disconnect all cables attached to the adapters and remove all of the adapters.
  3. Remove the removable media or disk drive enclosure assembly by pulling out the blue tabs at the bottom of the enclosure, then sliding the enclosure out approximately three centimeters.
  4. Remove and label the disk drives from the media or disk drive enclosure assembly.
  5. Remove one of the two memory DIMM pairs.
  6. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
  7. Turn on the power using either the management console or the white button.

  Does the managed system reach power on at hypervisor standby as indicated on the management console?

  No
  Go to PFW1548-7.

  Yes
  Go to the next step.

- **PFW1548-5**

  Were any memory DIMMs removed from system backplane?

  No
  Go to PFW1548-8.

  Yes
  Go to the next step.

- **PFW1548-6**

  1. Turn off the power, and remove the power cords.
  2. Replug the memory DIMMs that were removed from system backplane in PFW1548-2 in their original locations.
  3. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
  4. Turn on the power using either the management console or the white button.

  Does the managed system reach power on at hypervisor standby as indicated on the management console?

  No
  A memory DIMM in the pair you just replaced in the system is defective. Turn off the power, remove the power cords, and exchange the memory DIMM pair with new or previously removed memory DIMM pair. Repeat this step until the defective memory DIMM pair is identified, or both memory DIMM pairs have been exchanged.

  If your symptom did not change and both the memory DIMM pairs have been exchanged, call your service support person for assistance.

  If the symptom changed, check for loose cards and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

  Yes
  Go to the next step.

- **PFW1548-7**

  One of the FRUs remaining in the system unit is defective.

Note: If a memory DIMM is exchanged, ensure that the new memory DIMM is the same size and speed as the original memory DIMM.

1. Turn off the power, remove the power cords, and exchange the following FRUs, one at a time, in the order listed:

   a. Memory DIMMs. Exchange one pair at a time with new or previously removed DIMM pairs

   b. System backplane, location: U*n*-P1

   c. Power supplies, locations: U*n*-E1 and U*n*-E2.

2. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
3. Turn on the power using either the management console or the white button.

Does the managed system reach power on at hypervisor standby as indicated on the management console?

No

Reinstall the original FRU.

Repeat the FRU replacement steps until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

Yes

Go to Verifying a repair.

- **PFW1548-8**

1. Turn off the power.
2. Reconnect the system console.

   Notes:

   a. If an ASCII terminal has been defined as the firmware console, attach the ASCII terminal cable to the S1 connector on the rear of the system unit.

   b. If a display attached to a display adapter has been defined as the firmware console, install the display adapter and connect the display to the adapter. Plug the keyboard and mouse into the keyboard connector on the rear of the system unit.

3. Turn on the power using either the management console or the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.) After the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the diagnostic CD-ROM.
4. If the ASCII terminal or graphics display (including display adapter) is connected differently from the way it was previously, the console selection screen appears. Select a firmware console.
5. Immediately after the word *keyboard* is displayed, press the number 1 key on the directly attached keyboard, an ASCII terminal or management console. This activates the system management services (SMS).
6. Enter the appropriate password if you are prompted to do so.

Is the SMS screen displayed?

No

One of the FRUs remaining in the system unit is defective.

Exchange the FRUs that have not been exchanged, in the following order:

1. If you are using an ASCII terminal, go to the problem determination procedures for the display. If you do not find a problem, replace the system backplane at location U*n*-P1.
2. If you are using a graphics display, go to the problem determination procedures for the display. If you do not find a problem, complete the following steps:

   a. Replace the display adapter.

   b. Replace the backplane in which the graphics adapter is plugged.

      Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

      If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

      If the symptom changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

Yes
Go to the next step.

- **PFW1548-9**

1. Make sure the diagnostic CD-ROM is inserted into the optical drive.
2. Turn off the power and remove the power cords.
3. Use the cam levers to reconnect the disk drive enclosure assembly to the I/O backplane.
4. Reconnect the removable media or disk drive enclosure assembly by sliding the media enclosure toward the rear of the system, then pressing the blue tabs.
5. Plug in the power cords and wait for 01 in the upper-left corner of the operator panel display.
6. Turn on the power using either the management console or the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.) After the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the diagnostic CD-ROM.
7. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or an ASCII terminal keyboard.
8. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No
One of the FRUs remaining in the system unit is defective.

Exchange the FRUs that have not been exchanged, in the following order:

1. Optical drive
2. Removable media enclosure.
3. System backplane, U*n*-P1.

Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis procedures and follow the instructions for the new symptom.

Yes
Go to the next step.

- **PFW1548-10**

The system is working correctly with this configuration. One of the disk drives that you removed from the disk drive backplanes may be defective.

1. Make sure the diagnostic CD-ROM is inserted into the optical drive.
2. Turn off the power and remove the power cords.
3. Install a disk drive in the media or disk drive enclosure assembly.
4. Plug in the power cords and wait for the OK prompt to display on the operator panel display.
5. Turn on the power.
6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or an ASCII terminal keyboard.
7. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No

Exchange the FRUs that have not been exchanged, in the following order:

1. Last disk drive installed
2. Disk drive backplane.

Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis procedures and follow the instructions for the new symptom.

Yes

Repeat this step with all disk drives that were installed in the disk drive backplane.

After all of the disk drives have been reinstalled, go to the next step.

- **PFW1548-11**

The system is working correctly with this configuration. One of the devices that was disconnected from the system backplane may be defective.

1. Turn off the power and remove the power cords.
2. Attach a system backplane device (for example: system port 1, system port 2, USB, keyboard, mouse, Ethernet) that had been removed.

   After all of the device cables have been reattached, reattached the cables to the service processor one at a time.
3. Plug in the power cords and wait for 01 in the upper-left corner on the operator panel display.
4. Turn on the power using either the management console or the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.) After the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the diagnostic CD-ROM.
5. If the Console Selection screen is displayed, choose the system console.
6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
7. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No

The last device or cable that you attached is defective.

To test each FRU, exchange the FRUs in the following order:

1. Device and cable (last one attached)
2. System backplane, location: U*n*-P1.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis procedures and follow the instructions for the new symptom.

Yes

Repeat this step until all of the devices are attached. Go to the next step.

- **PFW1548-12**

  The system is working correctly with this configuration. One of the FRUs (adapters) that you removed may be defective.

  1. Turn off the power and remove the power cords.
  2. Install a FRU (adapter) and connect any cables and devices that were attached to the FRU.
  3. Plug in the power cords and wait for the OK prompt to display on the operator panel display.
  4. Turn on the power using either the management console or the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.) After the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the diagnostic CD-ROM.
  5. If the Console Selection screen is displayed, choose the system console.
  6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
  7. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No

  Go to the next step.

  Yes

  Repeat this step until all of the FRUs (adapters) are installed. Go to Verifying a repair.

- **PFW1548-13**

  The last FRU installed or one of its attached devices is probably defective.

  1. Make sure the diagnostic CD-ROM is inserted into the optical drive.
  2. Turn off the power and remove the power cords.
  3. Starting with the last installed adapter, disconnect one attached device and cable.
  4. Plug in the power cords and wait for the 01 in the upper-left corner on the operator panel display.
  5. Turn on the power using either the management console or the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.) After the system has reached hypervisor standby, activate a Linux or AIX partition by clicking the **Advanced** button on the Advanced activation screen. On the Advanced activation screen, select **Boot in service mode using the default boot list** to boot the diagnostic CD-ROM.
  6. If the Console Selection screen is displayed, choose the system console.
  7. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
  8. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No

  Repeat this step until the defective device or cable is identified or all devices and cables have been disconnected.

  If all the devices and cables have been removed, then one of the FRUs remaining in the system unit is defective.

To test each FRU, exchange the FRUs in the following order:

1. Adapter (last one installed)
2. System backplane, location: U*n*-P1.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

Yes

The last device or cable that you disconnected is defective. Exchange the defective device or cable then go to the next step.

- **PFW1548-14**

1. Follow the instructions on the screen to select the system console.
2. When the DIAGNOSTIC OPERATING INSTRUCTIONS screen is displayed, press Enter.
3. Select **Advanced Diagnostics Routines**.
4. If the terminal type has not been defined, you must use the option **Initialize Terminal** on the FUNCTION SELECTION menu to initialize the stand-alone diagnostic environment before you can continue with the diagnostics. This is a separate operation from selecting the console display.
5. If the NEW RESOURCE screen is displayed, select an option from the bottom of the screen.

   Note: Adapters and devices that require supplemental media are not shown in the new resource list. If the system has adapters or devices that require supplemental media, select option 1.
6. When the DIAGNOSTIC MODE SELECTION screen is displayed, press Enter.
7. Select **All Resources**. (If you were sent here from step PFW1548-18, select the adapter or device that was loaded from the supplemental media).

Did you get an SRN?

No

Go to step PFW1548-16.

Yes

Go to the next step.

- **PFW1548-15**

Look at the FRU part numbers associated with the SRN.

Have you exchanged all the FRUs that correspond to the failing function codes (FFCs)?

No

Exchange the FRU with the highest failure percentage that has not been changed.

Repeat this step until all the FRUs associated with the SRN have been exchanged or diagnostics run with no trouble found. Run diagnostics after each FRU is exchanged. Go to Verifying a repair.

Yes

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

- **PFW1548-16**

Does the system have adapters or devices that require supplemental media?

No

Go to step the next step.

Yes

Go to step PFW1548-18.

- **PFW1548-17**

Consult the PCI adapter configuration documentation for your operating system to verify that all adapters are configured correctly.

Go to Verifying a repair.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

- **PFW1548-18**

  1. Select **Task Selection**.
  2. Select **Process Supplemental Media** and follow the on-screen instructions to process the media. Supplemental media must be loaded and processed one at a time.

  Did the system return to the TASKS SELECTION SCREEN after the supplemental media was processed?

  No
  Go to the next step.

  Yes
  Press F3 to return to the FUNCTION SELECTION screen. Go to step PFW1548-14, substep "4" on page 96.

- **PFW1548-19**

  The adapter or device is probably defective.

  If the supplemental media is for an adapter, replace the FRUs in the following order:

  1. Adapter
  2. System backplane, location: U*n*-P1.

  If the supplemental media is for a device, replace the FRUs in the following order:

  1. Device and any associated cables
  2. The adapter to which the device is attached

  Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

  Go to Verifying a repair.

  This ends the procedure.

PFW1548: Memory and processor subsystem problem isolation procedure without a management console attached

This procedure is used to locate defective FRUs not found by normal diagnostics. For this procedure, diagnostics are run on a minimally configured system. If a failure is detected on the minimally configured system, the remaining FRUs are exchanged one at a time until the failing FRU is identified. If a failure is not detected, FRUs are added back until the failure occurs. The failure is then isolated to the failing FRU.

Perform the following procedure:

- **PFW1548-1**

  1. Ensure that the diagnostics and the operating system are shut down.

     Is the system at "service processor standby", indicated by 01 in the control panel?

     No
     Replace the system backplane, location: U*n*-P1. Return to step PFW1548-1.

     Yes
     Continue with substep "2" on page 98.

2. Turn on the power using either the white button or the ASMI menus.

   Does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the "Please define the System Console" screen displayed?

   No
   Go to PFW1548-3.

   Yes
   Go to PFW1548-2.

3. Insert the stand-alone diagnostic CD-ROM into the optical drive.

   Note: If you cannot insert the diagnostic CD-ROM, go to PFW1548-2.

4. When the word *keyboard* is displayed on an ASCII terminal or a directly attached keyboard, press the number 5 key.
5. If you are prompted to do so, enter the appropriate password.

   Is the "Please define the System Console" screen displayed?

   No
   Go to PFW1548-2.

   Yes
   Go to PFW1548-14.

- **PFW1548-2**

  1. Insert the stand-alone diagnostic CD-ROM into the optical drive.

     Note: If you cannot insert the diagnostic CD-ROM, go to step PFW1548-3.

  2. Turn on the power using either the white button or the ASMI menus. If the diagnostic CD-ROM is not in the optical drive, insert it now. If you are prompted to do so, enter the appropriate password.

  Is the "Please define the System Console" screen displayed?

  No
  Go to PFW1548-3.

  Yes
  Go to PFW1548-14.

- **PFW1548-3**

  1. Turn off the power.
  2. If you have not already done so, configure the service processor (using the ASMI menus) with the instructions in note “6” on page 82 at the beginning of this procedure, then return here and continue.
  3. Exit the service processor (ASMI) menus and remove the power cords.
  4. Disconnect all external cables (parallel, system port 1, system port 2, keyboard, mouse, USB devices, SPCN, Ethernet, and so on). Also disconnect all of the external cables attached to the service processor.

  Go to the next step.

- **PFW1548-4**

  1. If this is a deskside system, remove the service access cover. If this is a rack-mounted system, place the drawer into the service position and remove the service access cover. Also remove the front cover.
  2. Record the slot numbers of the PCI adapters and I/O expansion cards if present. Label and record the locations of all cables attached to the adapters. Disconnect all cables attached to the adapters and remove all of the adapters.
  3. Remove the removable media or disk drive enclosure assembly by pulling out the blue tabs at the bottom of the enclosure, then sliding the enclosure out approximately three centimeters.
  4. Remove and label the disk drives from the media or disk drive enclosure assembly.

5. Remove a memory DIMM pair.
6. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
7. Turn on the power using the white button.

Does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the "Please define the System Console" screen displayed?

No
Go to PFW1548-7.

Yes
Go to the next step.

- **PFW1548-5**

  Were any memory DIMMs removed from system backplane?

  No
  Go to PFW1548-8.

  Yes
  Go to the next step.

- **PFW1548-6**

  1. Turn off the power, and remove the power cords.
  2. Replug the memory DIMMs that were removed from the system backplane in PFW1548-2 in their original locations.
  3. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
  4. Turn on the power using the white button.

  Does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the "Please define the System Console" screen displayed?

  No
  A memory DIMM in the pair you just replaced in the system is defective. Turn off the power, remove the power cords, and exchange the memory DIMMs pair with new or previously removed memory DIMM pair. Repeat this step until the defective memory DIMM pair is identified, or both memory DIMM pairs have been exchanged.

  If your symptom did not change and both the memory DIMM pairs have been exchanged, call your service support person for assistance.

  If the symptom changed, check for loose cards and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

  Yes
  Go to the next step.

- **PFW1548-7**

  One of the FRUs remaining in the system unit is defective.

  Note: If a memory DIMM is exchanged, ensure that the new memory DIMM is the same size and speed as the original memory DIMM.

  1. Turn off the power, remove the power cords, and exchange the following FRUs, one at a time, in the order listed:
     a. Memory DIMMs. Exchange one pair at a time with new or previously removed DIMM pairs.
     b. System backplane, location: U*n*-P1
     c. Power supplies, locations: U*n*-E1 and U*n*-E2.
  2. Plug in the power cords and wait for 01 in the upper-left corner of the control panel display.
  3. Turn on the power using the white button.

Does the system reach an operating system login prompt, or if booting the stand-alone diagnostic CD-ROM, is the "Please define the System Console" screen displayed?

No

Reinstall the original FRU.

Repeat the FRU replacement steps until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

Yes

Go to Verifying a repair.

- **PFW1548-8**

1. Turn off the power.
2. Reconnect the system console.

   Notes:

   a. If an ASCII terminal has been defined as the firmware console, attach the ASCII terminal cable to the S1 connector on the rear of the system unit.

   b. If a display attached to a display adapter has been defined as the firmware console, install the display adapter and connect the display to the adapter. Plug the keyboard and mouse into the keyboard connector on the rear of the system unit.
3. Turn on the power using the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.)
4. If the ASCII terminal or graphics display (including display adapter) is connected differently from the way it was previously, the console selection screen appears. Select a firmware console.
5. Immediately after the word *keyboard* is displayed, press the number 1 key on the directly attached keyboard, or an ASCII terminal. This action activates the system management services (SMS).
6. Enter the appropriate password if you are prompted to do so.

Is the SMS screen displayed?

No

One of the FRUs remaining in the system unit is defective.

Exchange the FRUs that have not been exchanged, in the following order:

1. If you are using an ASCII terminal, go to the problem determination procedures for the display. If you do not find a problem, replace the system backplane at location U*n*-P1.
2. If you are using a graphics display, go to the problem determination procedures for the display. If you do not find a problem, complete the following steps:

   a. Replace the display adapter.

   b. Replace the backplane in which the graphics adapter is plugged.

      Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

      If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

      If the symptom changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

Yes

Go to the next step.

- **PFW1548-9**

  1. Make sure the diagnostic CD-ROM is inserted into the optical drive.
  2. Turn off the power and remove the power cords.
  3. Use the cam levers to reconnect the disk drive enclosure assembly to the I/O backplane.
  4. Reconnect the removable media or disk drive enclosure assembly by sliding the media enclosure toward the rear of the system, then pressing the blue tabs.
  5. Plug in the power cords and wait for 01 in the upper-left corner of the operator panel display.
  6. Turn on the power using the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.)
  7. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or an ASCII terminal keyboard.
  8. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No

  One of the FRUs remaining in the system unit is defective.

  Exchange the FRUs that have not been exchanged, in the following order:

  1. Optical drive
  2. Removable media enclosure.
  3. System backplane, U*n*-P1.

  Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis procedures and follow the instructions for the new symptom.

  Yes

  Go to the next step.

- **PFW1548-10**

  The system is working correctly with this configuration. One of the disk drives that you removed from the disk drive backplanes may be defective.

  1. Make sure the diagnostic CD-ROM is inserted into the optical drive.
  2. Turn off the power and remove the power cords.
  3. Install a disk drive in the media or disk drive enclosure assembly.
  4. Plug in the power cords and wait for the OK prompt to display on the operator panel display.
  5. Turn on the power.
  6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or an ASCII terminal keyboard.
  7. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No

  Exchange the FRUs that have not been exchanged, in the following order:

  1. Last disk drive installed
  2. Disk drive backplane.

  Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis procedures and follow the instructions for the new symptom.

Yes

Repeat this step with all disk drives that were installed in the disk drive backplane.

After all of the disk drives have been reinstalled, go to the next step.

- **PFW1548-11**

  The system is working correctly with this configuration. One of the devices that was disconnected from the system backplane may be defective.

  1. Turn off the power and remove the power cords.
  2. Attach a system backplane device (for example: system port 1, system port 2, USB, keyboard, mouse, Ethernet) that had been removed.

     After all of the device cables have been reattached, reattached the cables to the service processor one at a time.
  3. Plug in the power cords and wait for 01 in the upper-left corner on the operator panel display.
  4. Turn on the power using the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.)
  5. If the Console Selection screen is displayed, choose the system console.
  6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
  7. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No

  The last device or cable that you attached is defective.

  To test each FRU, exchange the FRUs in the following order:

  1. Device and cable (last one attached)
  2. System backplane, location: U*n*-P1.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to Problem Analysis procedures and follow the instructions for the new symptom.

  Yes

  Repeat this step until all of the devices are attached. Go to the next step.

- **PFW1548-12**

  The system is working correctly with this configuration. One of the FRUs (adapters) that you removed may be defective.

  1. Turn off the power and remove the power cords.
  2. Install a FRU (adapter) and connect any cables and devices that were attached to the FRU.
  3. Plug in the power cords and wait for the OK prompt to display on the operator panel display.
  4. Turn on the power using the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.)
  5. If the Console Selection screen is displayed, choose the system console.
  6. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.

7. Enter the appropriate password if you are prompted to do so.

Is the "Please define the System Console" screen displayed?

No
Go to the next step.

Yes
Repeat this step until all of the FRUs (adapters) are installed. Go to Verifying a repair.

- **PFW1548-13**

  The last FRU installed or one of its attached devices is probably defective.

  1. Make sure the diagnostic CD-ROM is inserted into the optical drive.
  2. Turn off the power and remove the power cords.
  3. Starting with the last installed adapter, disconnect one attached device and cable.
  4. Plug in the power cords and wait for the 01 in the upper-left corner on the operator panel display.
  5. Turn on the power using either the white button. (If the diagnostic CD-ROM is not in the optical drive, insert it now.)
  6. If the Console Selection screen is displayed, choose the system console.
  7. Immediately after the word *keyboard* is displayed, press the number 5 key on either the directly attached keyboard or on an ASCII terminal keyboard.
  8. Enter the appropriate password if you are prompted to do so.

  Is the "Please define the System Console" screen displayed?

  No
  Repeat this step until the defective device or cable is identified or all devices and cables have been disconnected.

  If all the devices and cables have been removed, then one of the FRUs remaining in the system unit is defective.

  To test each FRU, exchange the FRUs in the following order:

  1. Adapter (last one installed)
  2. System backplane, location: U*n*-P1.

  If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

  If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

  Yes
  The last device or cable that you disconnected is defective. Exchange the defective device or cable and then go to the next step.

- **PFW1548-14**

  1. Follow the instructions on the screen to select the system console.
  2. When the DIAGNOSTIC OPERATING INSTRUCTIONS screen is displayed, press Enter.
  3. Select **Advanced Diagnostics Routines**.
  4. If the terminal type has not been defined, you must use the option **Initialize Terminal** on the FUNCTION SELECTION menu to initialize the stand-alone diagnostic environment before you can continue with the diagnostics. This is a separate operation from selecting the console display.
  5. If the NEW RESOURCE screen is displayed, select an option from the bottom of the screen.

     Note: Adapters and devices that require supplemental media are not shown in the new resource list. If the system has adapters or devices that require supplemental media, select option 1.
  6. When the DIAGNOSTIC MODE SELECTION screen is displayed, press Enter.

7. Select **All Resources**. If you were sent here from step PFW1548-18, select the adapter or device that was loaded from the supplemental media.

Did you get an SRN?

No
Go to step PFW1548-16.

Yes
Go to the next step.

- **PFW1548-15**

Look at the FRU part numbers associated with the SRN.

Have you exchanged all the FRUs that correspond to the failing function codes (FFCs)?

No
Exchange the FRU with the highest failure percentage that has not been changed.

Repeat this step until all the FRUs associated with the SRN have been exchanged or diagnostics run with no trouble found. Run diagnostics after each FRU is exchanged. Go to Verifying a repair.

Yes
If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

- **PFW1548-16**

Does the system have adapters or devices that require supplemental media?

No
Go to step the next step.

Yes
Go to step PFW1548-18.

- **PFW1548-17**

Consult the PCI adapter configuration documentation for your operating system to verify that all adapters are configured correctly.

Go to Verifying a repair.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

- **PFW1548-18**

1. Select **Task Selection**.
2. Select **Process Supplemental Media** and follow the on-screen instructions to process the media. Supplemental media must be loaded and processed one at a time.

Did the system return to the TASKS SELECTION SCREEN after the supplemental media was processed?

No
Go to the next step.

Yes
Press F3 to return to the FUNCTION SELECTION screen. Go to step PFW1548-14, substep “4” on page 103.

- **PFW1548-19**

The adapter or device is probably defective.

If the supplemental media is for an adapter, replace the FRUs in the following order:

1. Adapter
2. System backplane, location: U*n*-P1.

If the supplemental media is for a device, replace the FRUs in the following order:

1. Device and any associated cables
2. The adapter to which the device is attached

Repeat this step until the defective FRU is identified or all the FRUs have been exchanged.

If the symptom did not change and all the FRUs have been exchanged, call service support for assistance.

If the symptom has changed, check for loose cards, cables, and obvious problems. If you do not find a problem, go to the Problem Analysis procedures and follow the instructions for the new symptom.

Go to Verifying a repair.

This ends the procedure.

Problems with noncritical resources

Use this procedure to help you determine the cause of problems with noncritical resources.

Procedure

1. Is there an SRC in an 8-character format available on the problem summary form?

   Note: If the operator has not filled out the problem summary form, go to the problem reporting procedure for the operating system in use.

   No: Continue with the next step.
   Yes: Perform problem analysis using the SRC. **This ends the procedure**.
2. Does the problem involve a workstation resource?
   - **No**: Continue with the next step.
   - **Yes**: Perform the following steps:
     - Check that the workstation is operational.
     - Verify that the cabling and addressing for the workstation is correct.
     - Perform any actions indicated in the system operator message.

   If you need further assistance, contact your next level of support. **This ends the procedure**.
3. Does the problem involve a removable media resource?

   No: Continue with the next step.
   Yes: Go to “Using the product activity log” on page 65 to resolve the problem. **This ends the procedure**.
4. Does the problem involve a communications resource?
   - **No**: Contact your next level of support. **This ends the procedure**.
   - **Yes**: Are there any system operator messages that indicate a communications-related problem has occurred?
     - **No**: Contact your next level of support. **This ends the procedure**.
     - **Yes**: Perform any actions indicated in the system operator message. If you need further assistance, contact your next level of support. **This ends the procedure**.

Intermittent problems

An intermittent problem is a problem that occurs for a short time, and then goes away.

About this task

The problem may not occur again until some time in the future, if at all. Intermittent problems cannot be made to appear again easily.

Some examples of intermittent problems are:

- A reference code appears on the control panel (the system attention light is on) but disappears when you power off, then power on the system. An entry does not appear in the Product Activity Log.
- An entry appears in the problem log when you use the Work with Problems (WRKPRB) command. For example, an expansion unit becomes powered off, but starts working again when you power it on.
- The workstation adapter is in a hang condition but starts working normally when it gets reset.

Note: You can get equipment for the following conditions from your branch office or installation planning representative:

- If you suspect that the air at the system site is too hot or too cold, you need a thermometer to check the temperature.
- If you suspect the moisture content of the air at the system site is too low or too high, use a wet/dry bulb to check the humidity. See "General intermittent problem checklist" on page 107 for more information.
- If you need to check AC receptacles for correct wiring, you need an ECOS tester, Model 1023-100, or equivalent tester. The tester lets you quickly check the receptacles. If you cannot find a tester, use an analog multimeter instead. Do not use a digital multimeter.

Follow the steps below to correct an intermittent problem:

Procedure

1. Read the information in "About intermittent problems" on page 106 before you attempt to correct an intermittent problem.

   Then, continue with the next step of this procedure.
2. Perform *all* steps in the "General intermittent problem checklist" on page 107.

   Then, continue with the next step of this procedure.
3. Did you correct the intermittent problem?

   Yes: This ends the procedure.

   No: Go to "Analyzing intermittent problems" on page 109. **This ends the procedure.**

About intermittent problems

An intermittent problem can show many different symptoms, so it might be difficult for you to determine the real cause without completely analyzing the failure.

To help with this analysis, you should determine as many symptoms as possible.

- The complete reference code is necessary to determine the exact failing area and the probable cause.
- Product activity log (PAL) information can provide time and device relationships.
- Information about environmental conditions when the failure occurred can be helpful (for example, an electrical storm occurring when the failure happened).

Note: If you suspect that an intermittent problem is occurring, increase the log sizes to the largest sizes possible. Select the PAL option on the Start a Service Tool display (see Using the product activity log for details).

Types of intermittent problems

Following are the major types of intermittent problems:

- Code (PTFs):
  - Licensed Internal Code
  - IBM i
  - Licensed program products
  - Other application software
- Configuration:

  - Non-supported hardware that is used on the system
  - Non-supported system configurations
  - Non-supported communication networks
  - Model and feature upgrades that are not performed correctly
  - Incorrectly configured or incorrectly cabled devices
- Environment:
  - Power line disturbance (for example, reduced voltage, a pulse, a surge, or total loss of voltage on the incoming AC voltage line)
  - Power line transient (for example, lightning strike)
  - Electrical noise (constant or intermittent)
  - Defective grounding or a ground potential difference
  - Mechanical vibration
- Intermittent hardware failure

General intermittent problem checklist

Use the following procedure to correct intermittent problems.

About this task

Performing these steps removes the known causes of most intermittent problems.

Procedure

1. Discuss the problem with the customer.

   Look for the following symptoms:

   - A reference code that goes away when you power off and then power on the system.
   - Repeated failure patterns that you cannot explain. For example, the problem occurs at the same time of day or on the same day of the week.
   - Failures that started after system relocation.
   - Failures that occurred during the time specific jobs or software were running.
   - Failures that started after recent service or customer actions, system upgrade, addition of I/O devices, new software, or program temporary fix (PTF) installation.
   - Failures occurring only during high system usage.
   - Failures occur when people are close to the system or machines are attached to the system.
2. Recommend that the customer install the latest cumulative PTF package, since code PTFs have corrected many problems that seem to be hardware failures.

   The customer can order the latest cumulative PTF package electronically through Electronic Customer Support or by calling the Software Support Center.
3. If you have not already done so, use the maintenance package to see the indicated actions for the symptom described by the customer.

   Attempt to perform the on-line problem analysis procedure first. If this is not possible, such as when the system is down, go to the Starting a repair action.

   Use additional diagnostic tools, if necessary, and attempt to recreate the problem.

   Note: Ensure that the service information you are using is at the same level as the operating system.
4. Check the site for the following environmental conditions:

   a) Any electrical noise that matches the start of the intermittent problems. Ask the customer such questions as:

      - Have any external changes or additions, such as building wiring, air conditioning, or elevators been made to the site?

    - Has any arc welding occurred in the area?
    - Has any heavy industrial equipment, such as cranes, been operating in the area?
    - Have there been any thunderstorms in the area?
    - Have the building lights become dim?
    - Has any equipment been relocated, especially computer equipment?

    If there was any electrical noise, find its source and prevent the noise from getting into the system.

   b) Site temperature and humidity conditions that are within the system specifications.

    See temperature and humidity design criteria in the Planning for the system topic relevant for your system.

   c) Poor air quality in the computer room:

    - Look for dust on top of objects. Dust particles in the air cause poor electrical connections and may cause disk unit failures.
    - Smell for unusual odors in the air. Some gases can corrode electrical connections.

   d) Any large vibration (caused by thunder, an earthquake, an explosion, or road construction) that occurred in the area at the time of the failure.

    Note: A failure that is caused by vibration is more probable if the server is on a raised floor.

5. Ensure that all ground connections are tight.

   These items reduce the effects of electrical noise. Check the ground connections by measuring the resistance between a conductive place on the frame to building ground or to earth ground. The resistance must be 1.0 ohm or less.

6. Ensure proper cable retention is used, as provided.

   If no retention is provided, the cable should be strapped to the frame to release tension on cable connections.

   Ensure that you pull the cable ties tight enough to fasten the cable to the frame bar tightly. A loose cable can be accidentally pulled with enough force to unseat the logic card in the frame to which the cable is attached. If the system is powered on, the logic card could be destroyed.

7. Ensure that all workstation and communications cables meet hardware specifications:

   - All connections are tight.
   - Any twinaxial cables that are not attached to devices must be removed.
   - The lengths and numbers of connections in the cables must be correct.
   - Ensure that lightning protection is installed on any twinaxial cables that enter or leave the building.

8. Perform the following:

   a) Review recent repair actions.

    Contact your next level of support for assistance.

   b) Review entries in the problem log (WRKPRB).

    Look for problems that were reported to the user.

   c) Review entries in the PAL, SAL, and service processor log. Look for a pattern:

    - SRCs on multiple adapters occurring at the same time
    - SRCs that have a common time-of-day or day-of-week pattern
    - Log is wrapping (hundreds of recent entries and no older entries)

    Check the PAL sizes and increase them if they are smaller than recommended.

   d) Review entries in the history log (`Display Log (DSPLOG)`).

    Look for a change that matches the start of the intermittent problems.

   e) Ensure that the latest engineering changes are installed on the system and on all system I/O devices.

9. Ensure that the hardware configuration is correct and that the model configuration rules have been followed.

   Use the **Display hardware configuration** service function (under SST or DST) to check for any missing or failed hardware.
10. Was a system upgrade, feature, or any other field bill of material or feature field bill of material installed just before the intermittent problems started occurring?

    No: Continue with the next step.

    Yes: Review the installation instructions to ensure that each step was performed correctly. Then, continue with the next step of this procedure.
11. Is the problem associated with a removable media storage device?

    No: Continue with the next step.

    Yes: Ensure that the customer is using the correct removable media storage device cleaning procedures and good storage media. Then, continue with the next step of this procedure.
12. Perform the following to help prevent intermittent thermal checks:

    - Ensure that the AMDs are working.
    - Exchange all air filters as recommended.
13. If necessary, review the intermittent problems with your next level of support and installation planning representative.

    Ensure that all installation planning checks were made on the system. Because external conditions are constantly changing, the site may need to be checked again. **This ends the procedure.**

Analyzing intermittent problems

This procedure enables you to begin analyzing an intermittent problem.

About this task

Use this procedure only after you have first reviewed the information in "About intermittent problems" on page 106 and gone through the "General intermittent problem checklist" on page 107.

Procedure

1. Is a reference code associated with the intermittent problem?

   No: Continue with the next step.

   Yes: Go to Reference codes. If the actions in the reference code tables do not correct the intermittent problem, return here and continue with the next step.
2. Is a symptom associated with the intermittent problem?

   No: Continue with the next step.

   Yes: Go to "Intermittent symptoms" on page 109. If the information there does not help to correct the intermittent problem, return here and continue with the next step.
3. Go to "Failing area intermittent isolation procedures" on page 110.

   If the information there does not help to correct the intermittent problem, return here and then, continue with the next step.
4. Send the data you have collected to your next level of support so that an Authorized Program Analysis Report (APAR) can be written.

   This ends the procedure.

Intermittent symptoms

Use the table below to find the symptom and description of the intermittent problem. Then perform the corresponding intermittent isolation procedures.

Although an isolation procedure may correct the intermittent problem, use your best judgment to determine if you should perform the remainder of the procedure shown for the symptom.

Note: If the symptom for the intermittent problem you have is not listed, go to "Failing area intermittent isolation procedures" on page 110.

Table 12. Intermittent symptoms

Symptom	Description	Isolation procedure
System powered off.	The system was operating correctly, then the system powered off. A 1xxx SRC may occur when this happens, and this SRC info should be logged in the service processor log.	INTIP09
System stops.	The system is powered on but is not operating correctly. No SRC is displayed. The system attention light is off and the processor activity lights may be on or off. Noise on the power-on reset line can cause the processor to stop.	INTIP18
System or subsystem runs slow.	The system or the subsystem is not processing at its normal speed.	INTIP20

Failing area intermittent isolation procedures

This procedure helps you determine how to resolve intermittent problems when you do not have a system reference code (SRC) or cannot determine the symptom.

About this task

Use this table only if you do not have a system reference code (SRC), or cannot find your symptom in "Intermittent symptoms" on page 109.

Procedure

1. Perform all of the steps in "General intermittent problem checklist" on page 107 for all failing areas.

   Then, continue with the next step.
2. Refer to the table below, and perform the following:
   a) Find the specific area of failure under **Failing area**.
   b) Look down the column of the area of failure until you find an X.
   c) Look across to the **Isolation procedure** column and perform the procedure indicated.
   d) If the isolation procedure does not correct the intermittent problem, continue down the column of the area of failure until you have performed all of the procedures shown for the failing area.
3. Although an isolation procedure may correct the intermittent problem, use your best judgment to determine if you should perform the remainder of the procedures shown for the failing area.

Results

Table 13. Failing area intermittent isolation procedures

Failing area						Isolation procedure to perform:
Power	**Work station I/O processor**	**Disk unit adapter**	**Comm-unication**	**Processor bus**	**Tape optical**	**Perform all steps in:**
X	X	X	X	X	X	"General intermittent problem checklist" on page 107
X	X			X		INTIP05

Table 13. Failing area intermittent isolation procedures (continued)

Failing area						Isolation procedure to perform:
Power	**Work station I/O processor**	**Disk unit adapter**	**Comm-unication**	**Processor bus**	**Tape optical**	**Perform all steps in:**
	X	X	X	X	X	INTIP07
X						INTIP09
X						INTIP14
		X				INTIP16
X	X	X	X	X	X	INTIP18
	X	X	X	X	X	INTIP20

IPL problems

Use these scenarios to help you diagnose your IPL problem.

Cannot perform IPL from the control panel (no SRC)

Use this procedure when you cannot perform an IBM i IPL from the control panel (no SRC).

About this task

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

Procedure

1. Perform the following:
   a) Verify that the power cable is plugged into the power outlet.
   b) Verify that power is available at the customer's power outlet.
2. Start an IPL by doing the following:
   a) Select Manual mode and IPL type A or B on the control panel. See Control panel functions for details.
   b) Power on the system. See Powering on and off.

   Does the IPL complete successfully?

   No: Continue with the next step.
   Yes: **This ends the procedure.**
3. Have all the units in the system become powered on that you expected to become powered on?

   Yes: Continue with the next step.
   No: Go to Power problems and find the symptom that matches the problem. **This ends the procedure.**
4. Is an SRC displayed on the control panel?

   - **Yes**: Go to Power problems and use the displayed SRC to correct the problem. **This ends the procedure.**

- **No**: For all models, exchange the following FRUs, one at a time. Refer to the remove and replace procedures for your specific system for additional information.

  a) SPCN card unit. See symbolic FRU TWRCARD.

  b) Power Supply. See symbolic FRU PWRSPLY. **This ends the procedure.**

Cannot perform IPL at a specified time (no SRC)

Use this procedure when you cannot perform an IBM i IPL at a specified time (no SRC). To correct the IPL problem, perform this procedure until you determine the problem and can perform an IPL at a specified time.

About this task

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

Procedure

1. Verify the following:

   a) The power cable is plugged into the power outlet.

   b) That power is available at the customer's power outlet.

2. Power on the system in normal mode. See Powering on and off.

   Does the IPL complete successfully?

   Yes: Continue with the next step.
   No: Go to the Starting a repair action procedure. **This ends the procedure.**

3. Have all the units in the system become powered on that you expected to become powered on?

   Yes: Continue with the next step.
   No: Go to Starting a repair action and find the symptom that matches the problem. **This ends the procedure.**

4. Verify the requested system IPL date and time by doing the following:

   a) On the command line, enter the Display System Value command:

   ```
   DSPSYSVAL QIPLDATTIM
   ```

   Observe the system value parameters.

   Note: The system value parameters are the date and time the system operator requested a timed IPL.

   ```
   +--------------------------------------------------------------------------------+
   |Display System Value                                                            |
   |System:  S0000000                                                               |
   |System value . . . . . . . . . . :  QIPLDATTIM                                  |
   |                                                                                |
   |Description  . . . . . . . . . . :  Date and time to automatically IPL         |
   |                                                                                |
   |                                                                                |
   |IPL date     . . . . . . . . . . :  MM/DD/YY                                    |
   |IPL time     . . . . . . . . . . :  HH:MM:SS                                    |
   +--------------------------------------------------------------------------------+
   ```

   Figure 1. Display for QIPLDATTIM

   b) Verify the system date. On the command line, enter the Display System Value command:

   ```
   DSPSYSVAL QDATE
   ```

   Check the system values for the date.

```
+--------------------------------------------------------------------------------+
|Display System Value                                                            |
|System:  S0000000                                                               |
|System value . . . . . . . . . :  QDATE                                         |
|                                                                                |
|Description  . . . . . . . . . :  System date                                   |
|                                                                                |
|Date          . . . . . . . . . :  MM/DD/YY                                     |
+--------------------------------------------------------------------------------+
```

Figure 2. Display for QDATE

Does the operating system have the correct date?

- **Yes**: Continue with this step.
- **No**: Set the correct date by doing the following:

  1) On the command line, enter the Change System Value command (`CHGSYSVAL QDATE VALUE('mmddyy')`).
  2) Set the date by entering

     mm=month
     dd=day
     yy=year

  3) Press **Enter**.

c) Verify the system time. On the command line, enter the Display System Value command: `DSPSYSVAL QTIME`

Check the system values for the time.

```
+--------------------------------------------------------------------------------+
|Display System Value                                                            |
|System:  S0000000                                                               |
|System value . . . . . . . . . :  QTIME                                         |
|                                                                                |
|Description  . . . . . . . . . :  Time of day                                   |
|                                                                                |
|Time         . . . . . . . . . :  HH:MM:SS                                      |
+--------------------------------------------------------------------------------+
```

Figure 3. Display for QTIME

Does the operating system have the correct time?

- **Yes**: Continue with this step.
- **No**: Set the correct time by doing the following:

  1) On the command line, enter the Change System Value command (`CHGSYSVAL QTIME VALUE('hhmmss')`).
  2) Set the time by entering

     hh=24 hour time clock
     mm=minutes
     ss=seconds

  3) Press **Enter** and then, continue with the next step.

5. Verify that the system can perform an IPL at a specified time by doing the following:

   a) Set the IPL time to 5 minutes past the present time by entering the Change System Value command (`CHGSYSVAL SYSVAL(QIPLDATTIM) VALUE('mmddyy hhmmss')`) on the command line.

   mm = month to power on
   dd = day to power on
   yy = year to power on
   hh = hour to power on

mm = minute to power on
ss = second to power on

b) Power off the system by entering the Power Down System Immediate command (`PWRDWNSYS *IMMED`) on the command line.

c) Wait 5 minutes.

Does the IPL start at the time you specified?

No: Continue with the next step.
Yes: **This ends the procedure.**

6. Power on the system in normal mode. See Powering on and off.

   Does the IPL complete successfully?

   Yes: Continue with the next step.
   No: Go to Starting a repair action. **This ends the procedure.**

7. Find an entry in the Service Action Log that matches the time, SRC, and/or resource that compares to the reported problem.

   a) On the command line, enter the Start System Service Tools command:

   ```
   STRSST
   ```

   If you cannot get to SST, select DST. See Dedicated service tools (DST) for details.

   Note: Do not IPL the system or partition to get to DST.

   b) On the Start Service Tools Sign On display, type in a user ID with service authority and password.

   c) Select **Start a Service Tool** > **Hardware Service Manager** > **Work with service action log**.

   d) On the Select Timeframe display, change the From: Date and Time to a date and time prior to when the customer reported having the problem.

   e) Find an entry that matches one or more conditions of the problem:

   - SRC
   - Resource
   - Time
   - FRU list (choose **Display the failing item information** to display the FRU list).

   Notes:

   a. All entries in the service action log represent problems that require a service action. It may be necessary to handle any problem in the log even if it does not match the original problem symptom.

   b. The information displayed in the date and time fields are the time and date for the first occurrence of the specific system reference code (SRC) for the resource displayed during the time range selected.

   Did you find an entry in the Service Action Log?

   No: Continue with the next step.
   Yes: Go to step "9" on page 115.

8. Exchange the following parts one at a time.

   See the remove and replace procedures for your specific system. After exchanging each part, return to step "5" on page 113 to verify that the system can perform an IPL at a specified time.

   Note: If you exchange the control panel or the system backplane, you must set the correct date and time by performing step "4" on page 112.

   ⚠ **Attention:** Before exchanging any part, power off the system. See Powering on and off.

   - System unit backplane (see symbolic FRU SYSBKPL)

- System control panel
- System control panel cable

Did the IPL complete successfully after you exchanged all of the parts listed above?

No: Contact your next level of support. **This ends the procedure.**
Yes: Continue with the next step.

9. Was the entry isolated (is there a Y in the Isolated column)?

   - **No**: Go to Reference codes and use the SRC indicated in the log. **This ends the procedure.**
   - **Yes**: Display the failing item information for the Service Action Log entry. Items at the top of the failing item list are more likely to fix the problem than items at the bottom of the list.

     Exchange the failing items one at a time until the problem is repaired. After exchanging each one of the items, verify that the item exchanged repaired the problem.

     Notes:

     a. For symbolic FRUs see Symbolic FRUs.
     b. When exchanging FRUs, refer to the remove and replace procedures for your specific system.
     c. After exchanging an item, go to Verifying the repair.

     After the problem has been resolved, close the log entry by selecting **Close a NEW entry** on the Service Actions Log Report display. **This ends the procedure.**

Cannot automatically perform an IPL after a power failure

Use this procedure when you cannot automatically perform an IBM i IPL after a power failure.

Procedure

1. Normal or Auto mode on the control panel must be selected when power is returned to the system.

   Is Normal or Auto mode on the control panel selected?

   Yes: Continue with the next step.
   No: Select **Normal** or **Auto** mode on the control panel. **This ends the procedure.**
2. Use the Display System Value command (DSPSYSVAL) to verify that the system value under QPWRRSTIPL on the Display System Value display is equal to 1.

   Is QPWRRSTIPL equal to 1?

   Yes: Contact your next level of support.
   No: Use the Change System Value command (CHGSYSVAL) to set QPWRRSTIPL equal to 1. **This ends the procedure.**

Power problems

Use the following table to learn how to begin analyzing a power problem.

Table 14. Analyzing power problems

Symptom	What you should do
The system unit does not power on.	See "Cannot power on system unit" on page 116.
The system unit does not power off.	See "Cannot power off system unit" on page 124.
The system does not remain powered on during a loss of incoming AC voltage and has an uninterruptible power supply (UPS) installed.	See the UPS user's guide that was provided with your unit.

Cannot power on system unit

Complete this procedure to correct the problem and power on the system.

About this task

For important safety information before you continue with this procedure, see “Power isolation procedures” on page 118.

Procedure

1. Attempt to power on the system. For more information about powering on your system, see Starting a system.

   Does the system power on, and is the system power status indicator light on continuously?

   Note: The system power status indicator flashes at the slower rate (one flash per two seconds) while powered off, and at the faster rate (one flash per second) during a normal power-on sequence.

   No: Continue with the next step.
   Yes: Go to step “13” on page 118.
2. Are there any characters displayed on the control panel (a scrolling dot might be visible as a character)?

   No: Continue with the next step.
   Yes: Go to step “5” on page 116.
3. Are the mainline AC power cables from the power supply, power distribution unit, or external uninterruptible power supply (UPS) to the customer's AC power outlet connected and seated correctly at both ends?

   Yes: Continue with the next step.
   No: Connect the mainline AC power cables correctly at both ends and go to step “1” on page 116.
4. Complete the following steps:
   a) Verify that the UPS is powered on (if it is installed).

      If the UPS will not power on, follow the service procedures for the UPS to ensure proper line voltage and UPS operation.

   b) Disconnect the mainline AC power cable or AC power jumper cable from the system's AC power connector at the system.
   c) Use a multimeter to measure the AC voltage at the system end of the mainline AC power cable or AC power jumper cable.

      Note: Some system models have more than one mainline AC power cable or AC power jumper cable. For these models, disconnect all the mainline AC power cables or AC power jumper cables and measure the AC voltage at each cable before you continue with the next step.

      Is the AC voltage 200 - 240 V AC, or 100 - 127 V AC?

      No: Go to step “8” on page 117.
      Yes: Continue with the next step.
5. Complete the following steps:
   a) Disconnect the mainline AC power cables from the power outlet.
   b) Exchange the system unit control panel and control panel cable (if present). See Part locations and location codes.
   c) Reconnect the mainline AC power cables to the power outlet.
   d) Attempt to power on the system.

   Does the system power on?

   No: Continue with the next step.
   Yes: The system unit control panel or control panel cable (if present) was the failing item. **This ends the procedure.**

6. Complete the following steps:
   a) Disconnect the mainline AC power cables from the power outlet.
   b) Exchange the power supply or supplies (U*n*-E1, U*n*-E2). See Part locations and location codes.
   c) Reconnect the mainline AC power cables to the power outlet.
   d) Attempt to power on the system. See Starting a system.

   Does the system power on?

   No: Continue with the next step.
   Yes: The power supply was the failing item. **This ends the procedure.**
7. Complete the following steps:
   a) Disconnect the mainline AC power cables.
   b) Replace the system backplane (U*n*-P1). See Part locations and location codes.
   c) Reconnect the mainline AC power cables to the power outlet.
   d) Attempt to power on the system.

   Does the system power on?

   No: Continue with the next step.
   Yes: The system backplane was the failing item. **This ends the procedure.**
8. Are you working with a system unit with a power distribution unit with tripped breakers?
   - **No:** Continue with the next step.
   - **Yes:** Complete the following steps:
     a. Reset the tripped power distribution breaker.
     b. Verify that the removable AC power cable is not the problem. Replace the cord if it is defective.
     c. If the breaker continues to trip, install a new power supply in each location until the defective one is found. **This ends the procedure.**
9. Does the system have an external UPS installed?

   Yes: Continue with the next step.
   No: Go to step “11” on page 117.
10. Use a multimeter to measure the AC voltage at the external UPS outlets. Is the AC voltage 200 - 240 V or 100 - 127 V AC?

    No: The UPS needs service. For 9910 type UPS, call IBM Service Support. For all other UPS types, have the customer call the UPS provider. In the meantime, go to step “12” on page 117 to bypass the UPS.
    Yes: Replace the AC power cable. See System parts for FRU part number. **This ends the procedure.**
11. Complete the following steps:
    a) Disconnect the mainline AC power cable from the customer's AC power outlet.
    b) Use a multimeter to measure the AC voltage at the customer's AC power outlet.

       Note: Some system models have more than one mainline AC power cable. For these models, disconnect all the mainline AC power cables and measure the AC voltage at all AC power outlets before you continue with this step.

       Is the AC voltage 200 - 240 V AC or 100 - 127 V AC?

       Yes: Exchange the mainline AC power cable. See System parts for the FRU part number. Then, go to step “1” on page 116.
       No: Inform the customer that the AC voltage at the power outlet is not correct. When the AC voltage at the power outlet is correct, reconnect the mainline AC power cables to the power outlet. **This ends the procedure.**
12. Complete the following steps to bypass the UPS unit:

a) Power off your system and the UPS unit.

b) Remove the signal cable that is used between the UPS and the system.

c) Remove any power jumper cords that are used between the UPS and the attached devices.

d) Remove the country or region-specific power cord that is used from the UPS to the wall outlet.

e) Use the correct power cord (the original country or region-specific power cord that was provided with your system) and connect it to the power inlet on the system. Plug the other end of this cord into a compatible wall outlet.

f) Attempt to power on the system.

Does the power-on standby sequence complete successfully?

Yes: Go to Verifying a repair. **This ends the procedure.**
No: Go to step "5" on page 116.

13. Display the selected IPL mode on the system unit control panel.

    Is the selected mode the same mode that the customer was using when the power-on failure occurred?

    No: Go to step "15" on page 118.
    Yes: Continue with the next step.

14. Is a function 11 reference code displayed on the system unit control panel?

    No: Go to step "16" on page 118.
    Yes: Return to Starting a repair action. **This ends the procedure.**

15. Complete the following steps:

    a) Power off the system. For more information about powering on and off your system, see Stopping a system.

    b) Select the mode on the system unit control panel that the customer was using when the power-on failure occurred.

    c) Attempt to power on the system.

    Does the system power on?

    Yes: Continue with the next step.
    No: Exchange the system unit control panel. See Part locations and location codes. **This ends the procedure.**

16. Continue the IPL. Does the IPL complete successfully?

    Yes: This ends the procedure.
    No: Return to Starting a repair action. **This ends the procedure.**

Power isolation procedures

Use power isolation procedures for isolating a problem in the power system. Use isolation procedures if there is not a management console attached to the server. If the server is connected to a management console, use the procedures that are available on the management console to continue FRU isolation.

Some field replaceable units (FRUs) can be replaced with the unit powered on. Follow the instructions in Part locations and location codes when directed to remove, exchange, or install a FRU.

The following safety notices apply throughout the power isolation procedures. Read all safety procedures before servicing the system and observe all safety procedures when performing a procedure.

⚠ **DANGER:** When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.

- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
  - For AC power, disconnect all power cords from their AC power source.
  - For racks with a DC power distribution panel (PDP), disconnect the customer’s DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected.
  - For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
  - For racks with a DC power distribution panel (PDP), connect the customer’s DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements.
- Do not continue with the inspection if any unsafe conditions are present.
- Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.

⚠ **DANGER:**

- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect:

  1. Turn off everything (unless instructed otherwise).
  2. For AC power, remove the power cords from the outlets.
  3. For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source.
  4. Remove the signal cables from the connectors.
  5. Remove all cables from the devices.

  To Connect:

  1. Turn off everything (unless instructed otherwise).
  2. Attach all cables to the devices.
  3. Attach the signal cables to the connectors.
  4. For AC power, attach the power cords to the outlets.
  5. For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP.
  6. Turn on the devices.

Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Cannot power on SPCN-controlled I/O expansion unit

You are here because an SPCN-controlled I/O expansion unit cannot be powered on, and might be displaying a 1*xxx*C62E reference code.

About this task

For important safety information before you continue with this procedure, see "Power isolation procedures" on page 118.

Procedure

1. Power on the system.
2. Start from SPCN 0 or SPCN 1 on the system unit. See Part locations and location codes, then go to the first unit in the SPCN frame-to-frame cable sequence that does not power on.

   Is the Data display background light on, or is the power-on LED flashing, or are there any characters displayed on the I/O expansion unit display panel?

   Note: The background light is a dim yellow light in the Data area of the display panel.

   Yes: Go to step "12" on page 122.
   No: Continue with the next step.
3. Use a multimeter to measure the AC voltage at the customer's AC power outlet.

   Is the AC voltage 200 - 240 V AC, or 100 - 127 V AC?

   - **Yes:** Continue with the next step.
   - **No:** Inform the customer that the AC voltage at the power outlet is not correct.

     This ends the procedure.
4. Is the mainline AC power cable from the AC module, power supply, or power distribution unit to the customer's AC power outlet connected and seated correctly at both ends?

   - **Yes:** Continue with the next step.
   - **No:** Connect the mainline AC power cable correctly at both ends.

     This ends the procedure.
5. Complete the following steps:

   a) Disconnect the mainline AC power cable from the AC module, power supply, or power distribution unit.

   b) Use a multimeter to measure the AC voltage at the AC module, power supply, or power distribution unit end of the mainline AC power cable.

   Is the AC voltage 200 - 240 V AC, or 100 - 127 V AC?

   No: Continue with the next step.
   Yes: Go to step "7" on page 121.
6. Are you working on a power distribution unit with tripped breakers?

   - **No:** Replace the mainline AC power cable or power distribution unit.

     This ends the procedure.
   - **Yes:** Complete the following steps:

     a. Reset the tripped power distribution breaker.
     b. Verify that the removable AC line cord is not the problem. Replace the cord if it is defective.
     c. Install a new power supply (one with the same part number as the one that is currently installed) in all power locations until the defective one is found.

This ends the procedure.

7. Does the unit that you are working on have AC power jumper cables installed?

   Note: The AC power jumper cables connect from the AC module, or the power distribution unit, to the power supplies.

   Yes: Continue with the next step.
   No: Go to step “11” on page 121.

8. Are the AC power jumper cables connected and seated correctly at both ends?

   - **Yes:** Continue with the next step.
   - **No:** Connect the AC power jumper cables correctly at both ends.

   This ends the procedure.

9. Complete the following steps:

   a) Disconnect the AC power jumper cables from the AC module or power distribution unit.

   b) Use a multimeter to measure the AC voltage at the AC module or power distribution unit (that goes to the power supplies).

   Is the AC voltage at the AC module or power distribution unit 200 - 240 V AC, or 100 - 127 V AC?

   - **Yes:** Continue with the next step.
   - **No:** Replace the following items (see System parts for location and part number information):
     - AC module
     - Power distribution unit

   This ends the procedure.

10. Complete the following steps:

    a) Connect the AC power jumper cables to the AC module, or power distribution unit.

    b) Disconnect the AC power jumper cable at the power supplies.

    c) Use a multimeter to measure the voltage input that the power jumper cables provide to the power supplies.

    Is the voltage 200 - 240 V AC or 100 - 127 V AC for each power jumper cable?

    - **No:** Exchange the power jumper cable.

    This ends the procedure.

    - **Yes:** Replace the following parts one at a time:
      a. I/O backplane
      b. Display unit
      c. Power supply 1
      d. Power supply 2
      e. Power supply 3

    This ends the procedure.

11. Complete the following steps:

    a) Disconnect the mainline AC power cable (to the expansion unit) from the customer's AC power outlet.

    b) Exchange the following FRUs, one at a time:

       - Power supply
       - I/O backplane

    c) Reconnect the mainline AC power cables (from the expansion unit) into the power outlet.

    d) Attempt to power on the system.

Does the expansion unit power on?

- **Yes:** The unit you exchanged was the failing item.

  This ends the procedure.

- **No:** Repeat this step and exchange the next FRU in the list. If you exchanged all of the FRUs in the list, ask your next level of support for assistance.

  This ends the procedure.

12. Is there a reference code displayed on the display panel for the I/O unit that does not power on?

    - **Yes:** Continue with the next step.
    - **No:** Replace the I/O backplane.

      This ends the procedure.

13. Is the reference code 1*xxxxx*2E?

    - **Yes:** Continue with the next step.
    - **No:** Use the new reference code and return to Start of call.

      This ends the procedure.

14. Do the SPCN optical cables (A) connect the failing unit (B) to the preceding unit in the chain or loop?

```
.---------.      (A) SPCN
| System  |  Optical Cables -.      .----- SPCN
| Unit    |                 |       V   Optical Adapter
| SPCN 0  |        .-.      V    .-.
'----.----'        | +------------+ |
     |             | +------------+ |
.----'----. .-------'-+    .-------'-+ .---------.
|  J15    | |Sec   J16|    |Sec   J15| |    Sec  |
|Sec Unit +-+UNIT  J15|    |Unit  J16+-+J15 Unit |
|  1      | |  2      |    |  3      | |      4  |
'---------' '---------'    '---------' '---------'
                               ^
                               |
                               '---- (B) Failing Unit
```

    Yes: Continue with the next step.

    No: Go to step “18” on page 123.

15. Remove the SPCN optical adapter (A) from the frame that precedes the frame that cannot become powered on.

```
.---------.          .--- (A) SPCN Optical Adapter
| System  |          |
| Unit    |          V
| SPCN 0  |         .-.           .-.
'----.----'         | +------------+ |
    322             | +------------+ |
.----'----. .-------'-+    .-------'-+ .---------.
|   J15   | |Sec   J16|    |Sec   J15| |    Sec  |
|Sec Unit +-+Unit  J15|    |Unit  J16+-+J15 Unit |
|  1      | |  2      |    |  3      | |      4  |
'---------' '---------'    '---------' '---------'
                                 ^
                                 |
                                 '-- Failing Unit
```

16. Complete the following steps:

    Notes:

    a. The cable might be connected to either J15 or J16.
    b. Use an insulated probe or jumper when you measure the voltage readings.

    a) Connect the negative lead of a multimeter to the system frame ground.

    b) Connect the positive lead of a multimeter to pin 2 of the connector from which you removed the SPCN optical adapter in the previous step of this procedure.

    c) Note the voltage reading on pin 2.

    d) Move the positive lead of the multimeter to pin 3 of the connector or SPCN card.

e) Note the voltage reading on pin 3.

Is the voltage on both pin 2 and pin 3 1.5 - 5.5 V DC?

- **Yes:** Continue with the next step.
- **No:** Exchange the I/O backplane.

  This ends the procedure.

17. Exchange the following FRUs, one at a time:

    a) In the failing unit (first frame with a failure indication), replace the I/O backplane.

    b) In the preceding unit in the string, replace the I/O backplane.

    c) SPCN optical adapter (A) in the preceding unit in the string.

    d) SPCN optical adapter (B) in the failing unit.

    e) SPCN optical cables (C) between the preceding unit in the string and the failing unit.

    This ends the procedure.

```
      (A) SPCN
       Optical            .----- (C) SPCN
       Adapter ----.      |      Optical Cables
                   |      |
                   |      |      .-- (B) SPCN
.---------.        |      |      |      Optical
| System  |        |      |      |      Adapter
| Unit    |        V      |      V
| SPCN 0  |      .-.      V     .-.
'----.----'      | +------------+ |
     |           | +------------+ |
.----'----. .-------'-+  .-------'-+ .---------.
|Sec J15  | |  Sec J16|  |Sec   J15| |    Sec  |
|Unit  J16+-+J15 Unit |  |Unit  J16+-+J15 Unit |
|  1      | |      2  |  |  3      | |      4  |
'---------' '---------'  '---------' '---------'
                               ^
                               |
                               |
                               '--- Failing Unit
```

18. Complete the following steps:

    a) Power off the system.

    b) Disconnect the SPCN frame-to-frame cable from the connector of the first unit that cannot be powered on.

    c) Connect the negative lead of a multimeter to the system frame ground.

    d) Connect the positive lead of the multimeter to pin 2 of the SPCN cable.

    Note: Use an insulated probe or jumper when you measure the voltage readings.

    e) Note the voltage reading on pin 2.

    f) Move the positive lead of the multimeter to pin 3 of the SPCN cable.

    g) Note the voltage reading on pin 3.

    Is the voltage on both pin 2 and pin 3 1.5 - 5.5 V DC?

    - **No:** Continue with the next step.
    - **Yes:** Exchange the following FRUs one at a time:

      a. In the failing unit, replace the I/O backplane.

      b. In the preceding unit in the string, replace the I/O backplane.

      c. SPCN frame-to-frame cable.

      This ends the procedure.

19. Complete the following steps:

    a) Follow the SPCN frame-to-frame cable back to the preceding unit in the string.

b) Disconnect the SPCN cable from the connector.

c) Connect the negative lead of a multimeter to the system frame ground.

d) Connect the positive lead of a multimeter to pin 2 of the connector.

Note: Use an insulated probe or jumper when you measure the voltage readings.

e) Note the voltage reading on pin 2.

f) Move the positive lead of the multimeter to pin 3 of the connector.

g) Note the voltage reading on pin 3.

Is the voltage on both pin 2 and pin 3 1.5 V - 5.5 V DC?

- **Yes:** Exchange the following FRUs one at a time:

  a. SPCN frame-to-frame cable.

  b. In the failing unit, replace the I/O backplane.

  c. In the preceding unit in the string, replace the I/O backplane.

  This ends the procedure.

- **No:** Exchange the I/O backplane from the unit from which you disconnected the SPCN cable in the previous step of this procedure.

This ends the procedure.

Cannot power off system unit

Use this procedure to analyze a failure of the normal command and control panel procedures to power off the system unit.

About this task

⚠ **Attention:** To prevent loss of data, ask the customer to verify that no interactive jobs are running before you complete this procedure.

For important safety information before you continue with this procedure, see "Power isolation procedures" on page 118.

Procedure

1. Attempt to power off the system.

   Does the system unit power off, and is the power indicator light flashing slowly?

   No: Continue with the next step.

   Yes: The system is not responding to normal power off procedures, which might indicate a Licensed Internal Code problem. Contact your next level of support. **This ends the procedure.**

2. Attempt to power off the system by using ASMI.

   Does the system power off?

   Yes: The system is not responding to normal power off procedures, which might indicate a Licensed Internal Code problem. Contact your next level of support. **This ends the procedure.**

   No: Continue with the next step.

3. Attempt to power off the system by using the control panel power button.

   Does the system power off?

   Yes: Continue with the next step.

   No: Go to step "5" on page 125.

4. Is there a reference code logged in the ASMI, the control panel, or the management console that indicates a power problem?

   Yes: Complete problem analysis for the reference code in the log. **This ends the procedure.**

   No: Contact your next level of support. **This ends the procedure.**

5. Ensure that no jobs are running on the system or partition, and verify that an uninterruptible power supply (UPS) is not powering the system. Then, continue with the next step.
6. Complete the following steps:
   a) Remove the system unit AC power cord from the external UPS or, if an external UPS is not installed, from the customer's AC power outlet.

      If the system unit has more than one AC line cord, disconnect all the AC line cords.

   b) Exchange the following FRUs one at a time. For more information about FRU locations and parts for the system that you are servicing, see Part locations and location codes and System parts.

      If the system unit is failing:

      1) Power supply. Go to step “7” on page 125.
      2) Service processor
      3) System control panel

      This ends the procedure.
7. A power supply might be the failing item.

   Attention: When you replace a redundant power supply, a 1*xxx*1504, 1*xxx*1514, 1*xxx*1524, or 1*xxx*1534 reference code might be logged in the error log. If you just removed and replaced the power supply in the location that is associated with this reference code, and the power supply came ready after the install, disregard this reference code. If you did not previously remove and replace a power supply, the power supply did not come ready after installation, or there are repeated fan fault errors after the power supply replacement, continue to follow these steps.

   Is the reference code 1*xxx*15*xx*?

   No: Continue with the next step.
   Yes: Complete the following steps:

   a) Find the unit reference code in one of the following tables to determine the failing power supply.
   b) Ensure that the power cables are properly connected and seated.
   c) Is the reference code 1*xxx*1500, 1*xxx*1510, 1*xxx*1520, or 1*xxx*1530 and is the failing unit configured with a redundant power supply option (or dual line cord feature)?
      - **Yes:** Complete “PWR1911” on page 126 before you replace parts.
      - **No:** Continue with step “7.d” on page 125.
   d) For more information about FRU locations for the system that you are servicing, see Part locations and location codes.
   e) Replace the failing power supply (see the following table to determine which power supply to replace).
   f) If the new power supply does not fix the problem, complete the following steps:
      1) Reinstall the original power supply.
      2) Try the new power supply in each of the other positions that are listed in the table.
      3) If the problem still is not fixed, reinstall the original power supply and go to the next FRU in the list.
      4) For reference codes 1*xxx*1500, 1*xxx*1510, 1*xxx*1520, and 1*xxx*1530, exchange the power distribution backplane if a problem persists after you replace the power supply.

Table 15. System unit

Unit reference code	Power supply
1510, 1511, 1512, 1513, 1514, 7110	E1
1520, 1521, 1522, 1523, 1524, 7120	E2

⚠ **Attention:** For reference codes 1500, 1510, 1520, and 1530, complete "PWR1911" on page 126 before you replace parts.

This ends the procedure.

8. Is the reference code 1*xxx*2600, 1*xxx*2603, 1*xxx*2605, or 1*xxx*2606?
   - **No:** Continue with the next step.
   - **Yes:** Complete the following steps:

   a) For more information about FRU locations for the system you are servicing, see Part locations and location codes.

   b) Replace the failing power supply.

   c) If the new power supply does not fix the problem, complete the following steps:

      1) Reinstall the original power supply.

      2) Try the new power supply in each of the other positions that are listed in the table.

      3) If the problem still is not fixed, reinstall the original power supply and go to the next FRU in the list.

      ⚠ **Attention:** Do not install power supplies P00 and P01 AC jumper cables on the same AC input module.

Table 16. Failing power supplies

System or feature code	**Failing power supply**
System unit	U*n*-E1, U*n*-E2

This ends the procedure.

9. Is the reference code 1*xxx*8455 or 1*xxx*8456?
   - **No:** Return to Starting a repair action. **This ends the procedure.**
   - **Yes:** One of the power supplies is missing, and must be installed. Use the following table to determine which power supply is missing, and install the power supply. For more information about FRU locations for the system you are servicing, see Part locations and location codes.

Table 17. Missing power supplies

Reference code	**Missing power supply**
1*xxx*8455	U*n*-E1
1*xxx*8456	U*n*-E2

This ends the procedure.

PWR1911
You are here because of a power problem on a system with more than one line cord. If the failing unit does not have more than one line cord, return to the procedure that sent you here or go to the next item in the FRU list.

About this task

The following steps are for the system unit, unless other instructions are given. For important safety information before you service the system, see "Power isolation procedures" on page 118.

Procedure

1. If an uninterruptible power supply is installed, verify that it is powered on before proceeding.
2. Are all the units powered on?

- **Yes:** Go to step “7” on page 128.
- **No:** On the unit that does not power on, complete the following steps:

  a. Disconnect the line cords from the unit that does not power on.

  b. Use a multimeter to measure the voltage at the system end of the line cords.

  Table 18. Correct voltage

Model or expansion drawer	**Correct AC voltage**	**Correct DC voltage**
5105-22E, 9008-22L, 9009-22A, 9009-22G, 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-22H, 9223-22S, 9223-42H, or 9223-42S	100 - 127 V or 200 - 240 V	-37.5 - -60 V or 192 - 400 V
9040-MR9	200 - 240 V	192 - 400 V
9080-M9S	200 - 240 V	192 - 400 V
EMX0 PCIe3 expansion drawer	100 - 127 V or 200 - 240 V	192 - 400 V

  c. Is the voltage correct (see Table 18 on page 127)?

  Yes: Continue with the next step.
  No: Go to step “5” on page 127.

3. Complete the following steps:

   a. Reconnect the line cords.

   b. Verify that the failing unit fails to power on.

   c. Replace the failing power supply. Use the following table to determine which power supply needs to be replaced, and then see Part locations and location codes for location, part number, and exchange procedure information.

   Table 19. Failing power supply for systems models and expansion drawers

Reference code	**System unit or expansion drawer**	**Failing item name**
1510	System unit	Power supply 1
	Expansion drawer	Power supply 1
1520	System unit	Power supply 2
	Expansion drawer	Power supply 2
1530	System unit	Power supply 3
1540	System unit	Power supply 4

   This ends the procedure.

4. Is the system a 9080-M9S?

   Yes: Continue with the next step.
   No: Go to step “6” on page 128.

5. Complete the following steps at the rear of the system:

   a) Disconnect the line cords from the connectors on the line cord conduit for the unit that does not power on.

   b) Use a multimeter to measure the voltage at the system end of the line cords.

c) Is the voltage correct (see Table 18 on page 127)?

Yes: Replace the line cord conduit. **This ends the procedure.**
No: Continue with the next step.

6. Complete the following steps:

   a) Disconnect the line cords from the customer's power outlet.

   b) Use a multimeter to measure the voltage at the customer's power outlet.

   Is the voltage correct (see Table 18 on page 127)?

   - **Yes:** Exchange the failing line cord.

     This ends the procedure.

   - **No:** Complete the following steps:

     a. Inform the customer that the voltage at the power outlet is not correct.

     b. Reconnect the line cords to the power outlet after the voltage at the power outlet is correct.

        This ends the procedure.

7. Is the reference code 1xxx00AC?

   - **No:** Continue with the next step.
   - **Yes:** This reference code might be caused by an outage. If the system will power on without an error, no parts need to be replaced.

     This ends the procedure.

8. Is the reference code 1xxx15x0?

   - **No:** Complete Problem Analysis by using the reference code.

     This ends the procedure.

   - **Yes:** Complete the following steps:

     a. Use the following table to locate the failing parts. For more information about locations, see Part locations and location codes.

Table 20. Power reference code table

System unit or expansion drawer	**Reference code**	**Locate these parts**
System unit	1xxx 1510	Power supply E1 and line cord 1
	1xxx 1520	Power supply E2 and line cord 2
	1xxx 1530	Power supply E3 and line cord 3
	1xxx 1540	Power supply E4 and line cord 4
Expansion drawer	1xxx 1510	Power supply 1 and line cord 1
	1xxx 1520	Power supply 2 and line cord 2

     b. Locate the line cord or the jumper cable for the reference code you are working on.

     c. Go to step “9” on page 128.

9. Complete the following steps:

   ⚠ **Attention:** Do not disconnect the other system line cords or the other jumper cables when powered on.

   a) For the reference code you are working on, disconnect either the jumper cable or the line cord from the power supply.

b) Use a multimeter to measure the voltage at the power supply end of the jumper cable **or** the line cord.

Is the voltage correct (see Table 18 on page 127)?

No: Continue with the next step.

Yes: Exchange the failing power supply. See Table 19 on page 127 for its position, and then see Part locations and location codes for part numbers and directions to the correct exchange procedures. **This ends the procedure.**

10. Complete the following steps:

a) Disconnect the line cords from the power outlet.

b) Use a multimeter to measure the voltage at the customer's power outlet.

Is the voltage correct (see Table 18 on page 127)?

- **Yes:** Exchange the following items, one at a time:
  - Failing line cord
  - Failing jumper cable (if installed)

  This ends the procedure.
- **No:** Complete the following steps:
  a. Inform the customer that the voltage at the power outlet is not correct.
  b. Reconnect the line cords to the power outlet after the voltage at the power outlet is correct.

  This ends the procedure.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled "Cookies, Web Beacons and Other Technologies".

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Electronic emission notices

Class A Notices

The following Class A statements apply to the IBM servers that contain the POWER9 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) “. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany

Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

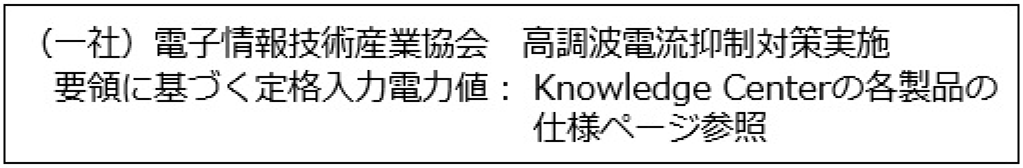

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

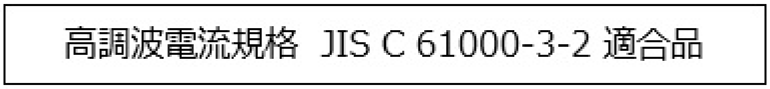

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

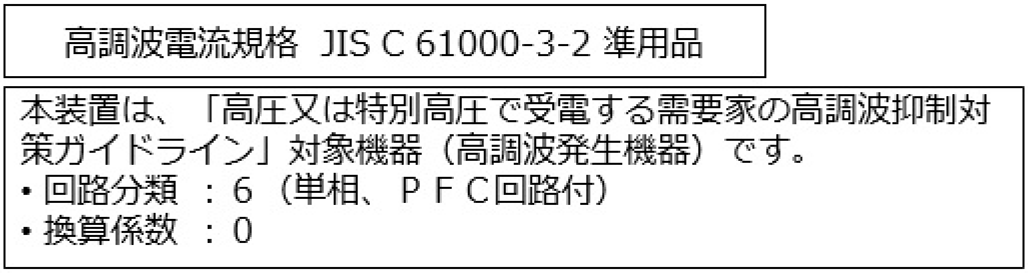

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

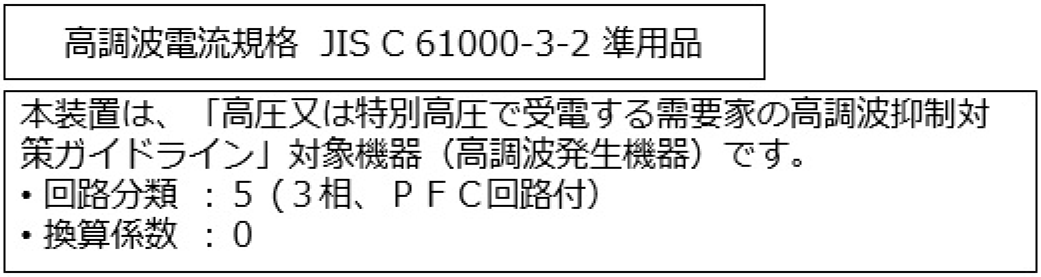

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

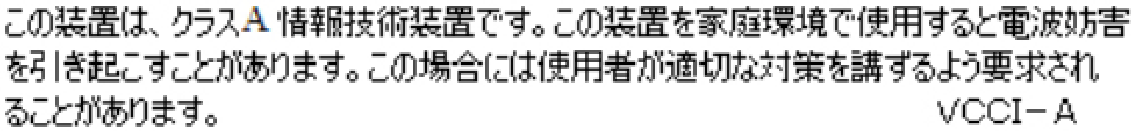

この装置は、クラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。　VCCI－A

Korea Notice

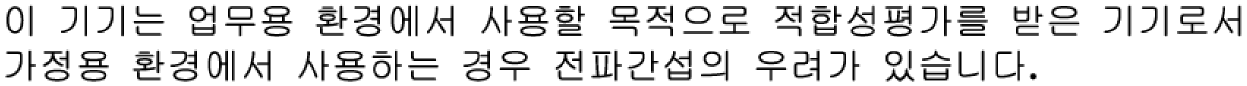

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

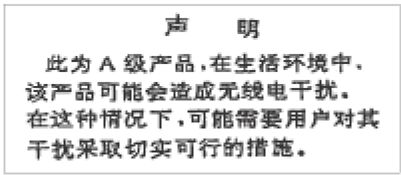

声　明

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰。
在这种情况下，可能需要用户对其
干扰采取切实可行的措施。

Russia Notice

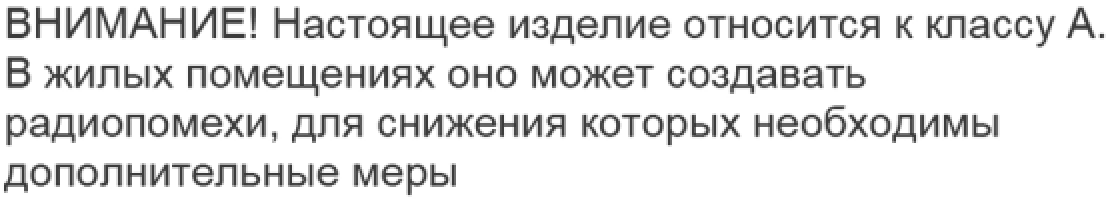

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

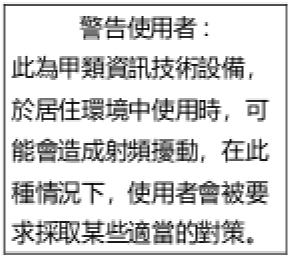

警告使用者：
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

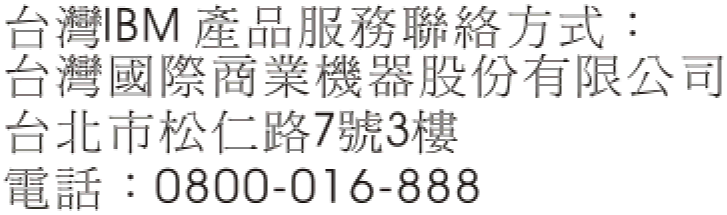

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) “. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany

Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

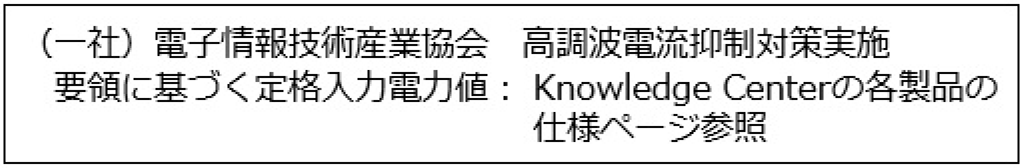

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値： Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

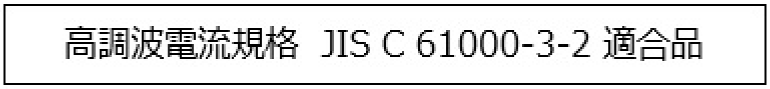

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

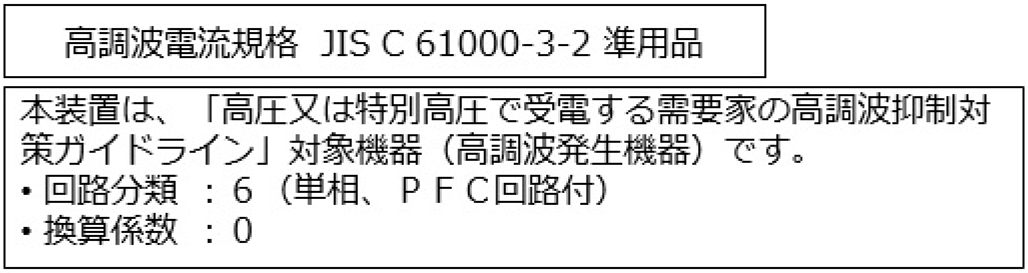

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

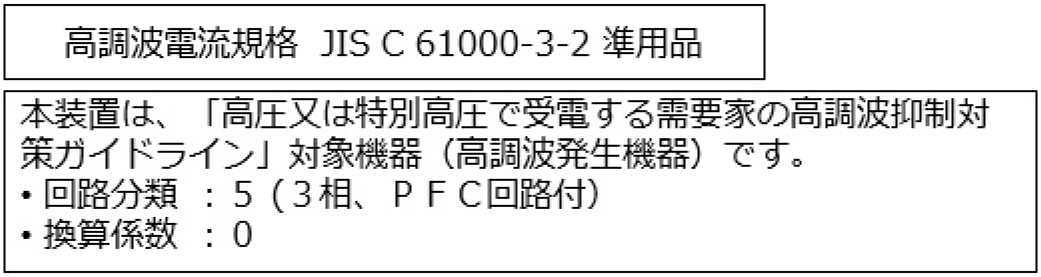

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

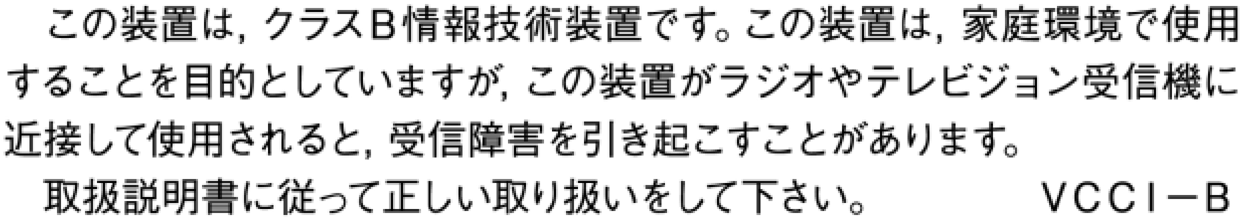

この装置は，クラスB情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。　VCCI－B

Taiwan Notice

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Installing and configuring the Hardware Management Console

IBM

Note

Before using this information and the product it supports, read the information in “Safety notices” on page v, “Notices” on page 135, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Hardware Management Console Version 9 Release 2 Maintenance Level 950 and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.

- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- 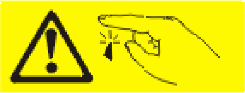 Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).

- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.

  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer’s DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

⚠ **CAUTION:**

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.

  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

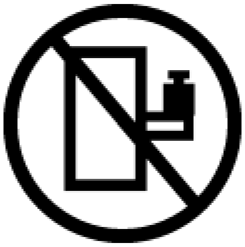

⚠ **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.
- Before extending the rack to the installation position, read the installation instructions.

- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

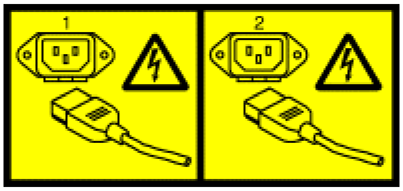

or

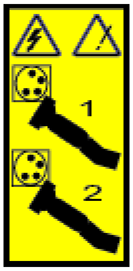

or

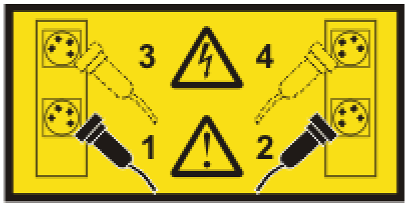

or

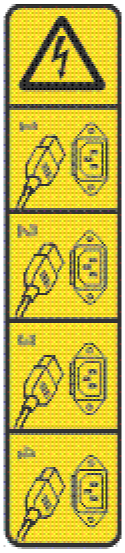

or

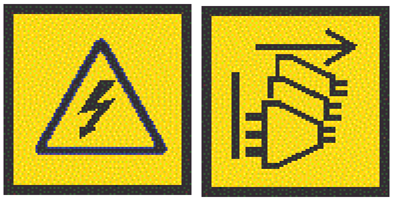

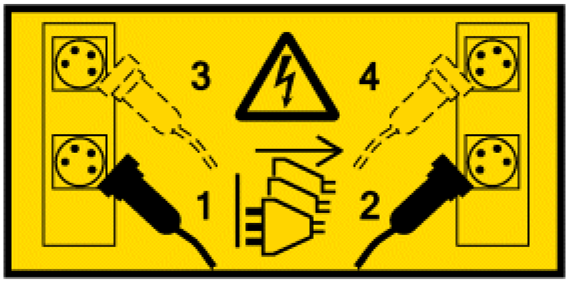

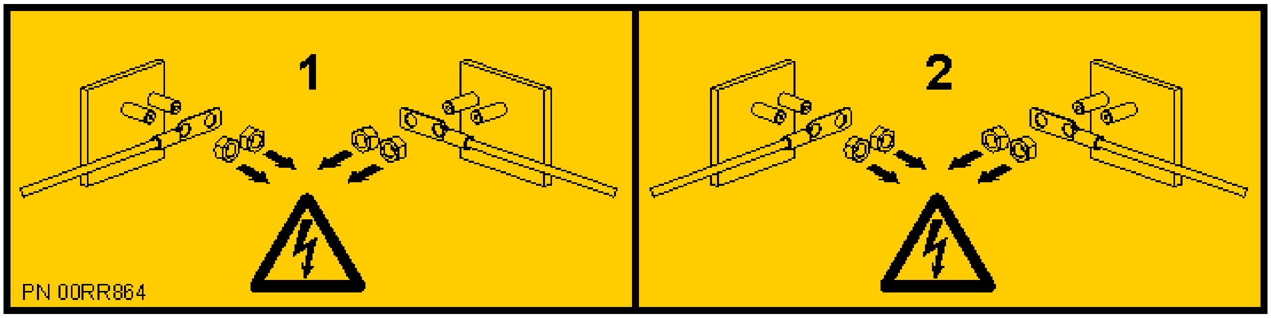

DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

CAUTION: A hot surface nearby. (L007)

(L008)

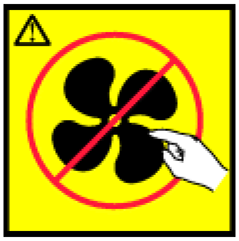

CAUTION: Hazardous moving parts nearby. (L008)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

CAUTION: This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

CAUTION: This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

CAUTION: Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

CAUTION: The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)

CAUTION: Regarding IBM provided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.
- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four (4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not

to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.

- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.
- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intrabuilding ports of this equipment are suitable for connection to intrabuilding or unexposed wiring or cabling only. The intrabuilding ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The ac-powered system does not require the use of an external surge protection device (SPD).

The dc-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The dc-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Installing and configuring the Hardware Management Console

Learn how to install the Hardware Management Console (HMC) hardware, connect it to your managed system, and configure it for use. You can perform these tasks yourself, or contact a service provider to perform these tasks for you. You might be charged a fee by the service provider for this service.

What's new in Installing and configuring the HMC

Read about new or significantly changed information in the Installing and configuring the HMC topic since the previous update of the topic collection.

April 2021

- Added the following topics:
  - "Installing the IBM Power Systems HMC (7063-CR2) into a rack" on page 4
  - "Prerequisites for installing the rack-mounted 7063-CR2 system" on page 4
  - "Completing inventory for your system" on page 5
  - "Determining and marking the location in the rack for the 7063-CR2 system" on page 5
  - "Attaching the adjustable rails to the system chassis and to the rack" on page 7
  - "Attaching the fixed rails to the system chassis and to the rack" on page 8
  - "Installing the system into the rack and connecting and routing power cables" on page 10
  - "Cabling the rack-mounted 7063-CR2 HMC" on page 10
  - "Configuring the 7063-CR2 HMC" on page 12

November 2020

- Updated the following topics:
  - "Installation and configuration tasks" on page 2
  - "Installing the 7042-CR9 HMC into a rack" on page 23
  - "Securing the HMC" on page 123
  - "HMC port locations" on page 130

July 2020

- Updated the following topics:
  - "Installing the HMC virtual appliance " on page 38
  - "HMC port locations" on page 130

October 2019

- Updated the following topics:
  - "Installing the HMC virtual appliance " on page 38
  - "Securing the HMC" on page 123

February 2019

- Added the following topics:

  - “Securing the HMC” on page 123
  - “Enhanced password policy” on page 125
  - “Solving common problems while securing the HMC” on page 126
  - “Security profiles: Global Data Protection Regulation (GDPR) and Payment Card Industry Data Security Standard (PCI-DSS) ” on page 128

August 2018

- Updated the following topics:
  - “Configuring the 7063-CR1 HMC” on page 21
  - “HMC port locations” on page 130

December 2017

- Added information for IBM Power Systems servers that contain the POWER9 processor.

Installation and configuration tasks

Learn about the tasks that are associated with different HMC installation and configuration tasks.

Learn about, at a high level, the tasks you must complete when you install and configure your HMC. You can install and configure your HMC in different ways. Find the situation that best matches the task that you want to complete.

Notes:

- If you are managing POWER9™ processor-based servers, the HMC must be at Version 9.1.0, or later. For more information, see “Determining your HMC machine code version and release” on page 113.
- Hardware Management Console Version 9.2.950, or later is not supported on the HMC 7042 Machine Type. For more information about the HMC versions for your 7042 HMC, see the HMC release notes that is available in the Fix Central website.

Installing and configuring a new HMC with a new server

Learn more about the high-level tasks you must complete when you install and configure a new HMC with a new server.

Table 1. Tasks that you need to complete when you install and configure a new HMC with a new server

Task	Where to find related information
1. Gather information and complete the Preinstallation Configuration worksheet.	“Preinstallation configuration worksheet for the HMC” on page 60 “Preparing for HMC configuration” on page 59
2. Unpack the hardware.	
3. Cable the HMC hardware.	For 7063-CR2 HMC:“Cabling the rack-mounted 7063-CR2 HMC” on page 10 For 7063-CR1 HMC: “Cabling the rack-mounted 7063-CR1 HMC” on page 19
4. Power on the HMC by pressing the power button.	
5. Log in and start the HMC web application.	
6. Use the HMC menus to configure the HMC.	“Configuring the HMC by using the menus ” on page 67

Table 1. Tasks that you need to complete when you install and configure a new HMC with a new server (continued)

Task	Where to find related information
7. Attach the server to the HMC.	

Updating and upgrading your HMC code

Learn more about the high-level tasks you must complete when you update and upgrade your HMC code.

If you have an existing HMC and want to update or upgrade your HMC code, you must complete the following high-level tasks:

Table 2. Tasks that you need to complete when you update or upgrade HMC code

Task	Where to find related information
1. Obtain the upgrade.	“Upgrading your HMC software” on page 118
2. View the existing HMC machine code level.	
3. Back up the managed system’s profile data.	
4. Back up HMC data.	
5. Record the current HMC configuration information.	
6. Record remote command status.	
7. Save upgrade data.	
8. Upgrade the HMC software.	
9. Verify that the HMC machine code upgrade installed successfully	

Adding a second HMC to an existing installation

Learn more about the high-level tasks you must complete when you add a second HMC to your managed system.

If you have an existing HMC and managed system and want to add a second HMC to this configuration, complete the following steps:

Table 3. Tasks that you need to complete when you add a second HMC to an existing installation

Task	Where to find related information
1. Ensure that your HMC hardware supports HMC Version 7 code.	
2. Gather information and complete the Preinstallation Configuration worksheet.	“Preinstallation configuration worksheet for the HMC” on page 60
3. Unpack the hardware.	
4. Cable the HMC hardware.	For 7063-CR2 HMC:“Cabling the rack-mounted 7063-CR2 HMC” on page 10 For 7063-CR1 HMC: “Cabling the rack-mounted 7063-CR1 HMC” on page 19
5. Power on the HMC by pressing the power button.	

Table 3. Tasks that you need to complete when you add a second HMC to an existing installation (continued)

Task	Where to find related information
6. Log in to the HMC.	
7. The HMC code levels must match. Change the code on one of the HMCs to match the code on the other.	"Determining your HMC machine code version and release" on page 113 "Upgrading your HMC software" on page 118
8. Access the Guided setup wizard or use the HMC menus to configure the HMC.	"Configuring the HMC by using the menus " on page 67
9. Configure this HMC for service by using the Call-Home Setup wizard.	"Configuring the HMC so that it can connect to service and support by using the call-home setup wizard" on page 106
10. Attach the server to the HMC.	

Setting up the HMC

You must set up the HMC hardware before you configure the HMC software. Learn more about setting up a desk-side HMC or a rack-mounted HMC.

Installing the IBM Power Systems HMC (7063-CR2) into a rack

Learn how to install the IBM Power Systems HMC (7063-CR2) into a rack.

You can view the online installation documentation, or you can print the PDF version of the same information. To view or print the PDF version, see Installing and configuring the Hardware Management Console.

Prerequisites for installing the rack-mounted 7063-CR2 system

Use the information to understand the prerequisites that are required for installing the system.

About this task

⚠ **CAUTION:** This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

You might need to read the following documents before you begin to install the server:

- The latest version of this document is maintained online, see Installing the 7063-CR2 into a rack (http://www.ibm.com/support/knowledgecenter/POWER9/p9hai/p9hai_install7063cr2_kickoff.htm).
- To plan your server installation, see Site and hardware planning.

Procedure

1. Ensure that you have the following items before starting your installation:
   - Size 2 Phillips screwdriver
   - Flat-head screwdriver
   - T25 screwdriver
   - Box cutter
   - Electrostatic discharge (ESD) wrist strap
   - Rack with one Electronic Industries Association (EIA) unit (1U) of space.

Notes:

- If you do not have a rack that is installed, install the rack. For instructions, see Racks and rack features (http://www.ibm.com/support/knowledgecenter/POWER9/p9hbf/p9hbf_9xx_kickoff.htm).
- The power supply ratings are 100 to 127 V ac, 9 A (x2), 200 to 240 V ac, 4.5 A (x2); 50 or 60 Hz.

2. Remove the shipping brackets on the system.
3. Continue with "Completing inventory for your system" on page 5.

Completing inventory for your system

Use this information to complete inventory for your system.

Procedure

1. Verify that you received all the boxes you ordered.
2. Unpack the server components as needed.
3. Complete a parts inventory and verify that you have received all the parts that you ordered before you install each server component.

   Note:

   Your order information is included with your product. You can also obtain order information from your marketing representative or the IBM Business Partner.

   If you have incorrect, missing, or damaged parts, consult any of the following resources:

   - Your IBM reseller.
   - IBM Rochester manufacturing automated information line at 1-800-300-8751 (United States only).
   - The Directory of worldwide contacts website (http://www.ibm.com/planetwide). Select your location to view the service and support contact information.
4. Continue with "Determining and marking the location in the rack for the 7063-CR2 system" on page 5.

Determining and marking the location in the rack for the 7063-CR2 system

You need to determine where to install the system unit into the rack.

Procedure

1. Read the Rack safety notices (http://www.ibm.com/support/knowledgecenter/POWER9/p9hbf/p9hbf_racksafety.htm).
2. Determine where to place the system unit in the rack. As you plan for installing the system unit in a rack, consider the following information:
   - Organize larger and heavier units into the lower part of the rack.
   - Plan to install system units into the lower part of the rack first.
   - Record the Electronic Industries Alliance (EIA) locations in your plan.
3. If necessary, remove the filler panels to allow access to the inside of the rack enclosure where you plan to place the unit, as shown in Figure 1 on page 6.

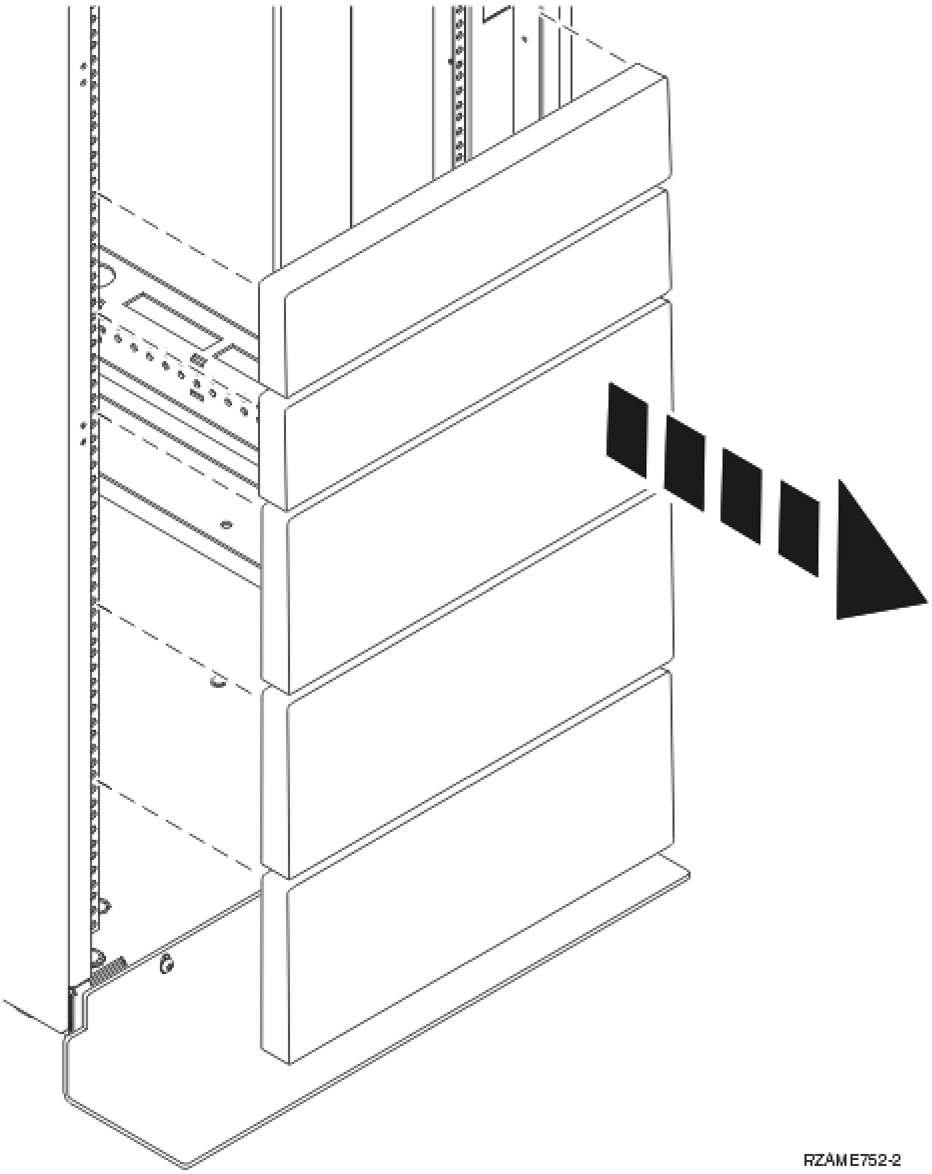

Figure 1. Removing the filler panels

4. Determine where to place the system in the rack. Record the EIA location.
5. Facing the front of the rack and working from the right side, use tape, a marker, or pencil to mark the lower hole of each EIA unit.
6. Repeat step “5” on page 6 for the corresponding holes located on the left side of the rack.
7. Go to the rear of the rack.
8. On the right side, find the EIA unit that corresponds to the bottom EIA unit marked on the front of the rack.
9. Mark the bottom EIA unit.
10. Mark the corresponding holes on the left side of the rack.
11. Continue with “Attaching the adjustable rails to the system chassis and to the rack” on page 7 to attach the adjustable rails or continue with “Attaching the fixed rails to the system chassis and to the rack” on page 8 to attach the fixed rails.

Attaching the adjustable rails to the system chassis and to the rack

You must install the rails onto the chassis and into the rack. Use this procedure to perform this task.

About this task

⚠ **Attention:** To avoid rail failure and potential danger to yourself and to the unit, ensure that you have the correct rails and fittings for your rack. If your rack has square support flange holes or screw-thread support flange holes, ensure that the rails and fittings match the support flange holes that are used on your rack. Do not install mismatched hardware by using washers or spacers. If you do not have the correct rails and fittings for your rack, contact your IBM reseller.

Note: The 1 EIA units in racks are measured in vertical increments of 44.45 mm (1.75 in.) each. Each 44.45 mm (1.75 in.) increment is called an “EIA.” In some countries, the same increment can be referred to as a “U.”

Note: The system requires 1 EIA rack unit (1U) of space.

Ensure that you have the necessary parts to install the rails. The following parts are included with the rail kit:

- 4 - Philips 6.35 mm (0.25 in.) screws
- 2 - rack and slide bracket rail assemblies
- 2 - HMC slide brackets
- 10 - nut clips for square EIA mounting holes
- 10 - nut clips for round EIA mounting holes
- 10 - M5 hex flange screws

Procedure

1. Remove the rail pieces from the packaging and put them on a work surface.
2. Identify 1U space in the rack of the HMC.
3. To attach the slide brackets to the HMC, perform the following tasks:
   a. Identify the right slide bracket.
   b. Align the holes on the right slide bracket with the slide bracket pins located on the right side of the HMC. Ensure that all the pins are aligned with the bracket holes.
   c. Push the HMC slide bracket toward the rear side of the HMC until it is fully locked into position.
   d. Secure the right slide bracket to the right side of the HMC workstation by installing two Philips 6.35 mm (0.25 in.) screws into the screw holes.
   e. Repeat the steps “3.a” on page 7 - “3.d” on page 7 to install the left slide bracket to the left side of the HMC workstation.
4. Move to the front of the rack.
   a. On the left side, install three nut clips into the three holes on the front edge of the rack in the 1U slot that is designated for the HMC.

      Note: The rail kit includes nut clips for both square and round rack holes. Ensure that you use appropriate nut clips that match the holes in the rack.
   b. Repeat step “4.a” on page 7 on the right side of the rack.
5. Move to the rear of the rack.
   a. On the left side, install two nut clips into the upper and lower holes on the front edge of the rack in the 1U slot that is designated for the HMC.

      Note: The middle hole must remain empty.
   b. Repeat step “5.a” on page 7 on the right side of the rack.

6. To install the HMC slide rails into the rack, perform the following steps:

    a. Measure the depth of the rack. The depth must be between 558.8 mm (22 in.) and 863.6 mm (34 in.).

    b. Place the HMC slide rails on a flat surface and locate the preinstalled screws.

       Note: The slide rails have four screw holes.

    c. Loosen the preinstalled screws on the slide rails enough that the rails can be moved in and out easily.

    d. Based on the depth of the rack measured in step "6.a" on page 8, you must adjust the screws on the rails.

       i) If the depth of the rack is between 558.8 mm (22 in.) and 698.5 mm (27.5 in.), attach the screws to the first and the third holes.

       ii) If the depth of the rack is between 698.5 mm (27.5 in.) and 863.6 mm (34 in.), attach the screws to the second and the fourth holes.

       Notes:

       - The first hole is always the hole closest to the end of the slide rail. The third and fourth holes are located close together.
       - Ensure that the screws are loose enough so that the length of the slide rail can be slightly adjusted while it is being installed in the rack.

7. At the front of the rack, install the HMC slide rails into the rack by performing the following steps:

    a. Locate the left slide rail assembly.

    b. Orient the rail assembly so that the end with the closest screw hole (the first hole) goes into the rack first. Ensure that the screw heads face the inside of the rack. The open slot of the rail assembly is closest to the front of the rack.

    c. On the left side of the rack, connect the flange on the end of the slide rail to the front edge of the rack by using two M5 screws, leaving the middle hole open. Ensure that the rail assembly is left slightly loose on the front of the rack to allow for the HMC to be inserted.

8. At the rear of the rack, on the right side, pull the free end of the slide rail toward the rear and secure the flange of the slide rail to the rack by using two M5 screws, leaving the middle screw hole open.

9. Repeat the step "7" on page 8 and step "8" on page 8 to install the right slide rail assembly on the right side of the rack.

10. In front of the rack, install the HMC workstation into the rack by performing the following steps:

    a. Holding the HMC workstation level, insert the slide brackets into the HMC slide rails that you installed in the previous step. Push the HMC forward until the flanges on the front of the HMC are flush with the open screw holes on the front of the rack.

    b. Connect the HMC to the left side of the frame using one M5 screw. Repeat this step on the right side of the rack.

11. Continue with "Installing the system into the rack and connecting and routing power cables" on page 10.

Attaching the fixed rails to the system chassis and to the rack

You must install the rails onto the chassis and into the rack. Use this procedure to perform this task.

About this task

⚠ **Attention:** To avoid rail failure and potential danger to yourself and to the unit, ensure that you have the correct rails and fittings for your rack. If your rack has square support flange holes or screw-thread support flange holes, ensure that the rails and fittings match the support flange holes that are used on your rack. Do not install mismatched hardware by using washers or spacers. If you do not have the correct rails and fittings for your rack, contact your IBM reseller.

Note: The 1 EIA unit in racks are measured in vertical increments of 44.45 mm (1.75 in.) each. Each 44.45 mm (1.75 in.) increment is called an “EIA.” In some countries, the same increment can be referred to as a “U.”

Note: The system requires 1 EIA rack unit (1U) of space.

Ensure that you have the necessary parts to install the rails. The following parts are included with the rail kit:

- 4 - Philips 6.35 mm (0.25 in.) screws
- 2 - Inner rails
- 2 - HMC support rails
- 2 - Nut clips for square EIA mounting holes
- 2 - Nut clips for round EIA mounting holes
- 8 - M5 hex flange screws

Procedure

1. Remove the rail pieces from the packaging and put them on a work surface.
2. Identify 1U space in the rack of the HMC.
3. To attach the inner rails to the HMC, perform the following tasks:
    a. Identify the right inner rail.
    b. Align the holes on the right inner rail with the inner rail pins located on the right side of the HMC. Ensure that all the pins are aligned with the inner rail holes.
    c. Push the HMC inner rail toward the front side of the HMC until it is fully locked into position.
    d. Secure the right inner rail to the right side of the HMC workstation by installing two Philips 6.35 mm (0.25 in.) screws into the screw holes.
    e. Repeat the steps 3.a - “3.d” on page 9 to install the left inner rail to the left side of the HMC workstation.
4. Move to the front of the rack. On the left side, install one nut clips into the hole on the front edge of the rack in the 1U slot that is designated for the HMC.

    Note: The rail kit includes nut clips for both square and round rack holes. Ensure that you use appropriate nut clips that match the holes in the rack.
5. Move to the rear of the rack. On the left side, install one nut clips into the middle hole on the front edge of the rack in the 1U slot that is designated for the HMC.
6. At the front of the rack, install the HMC support rails into the rack by performing the following steps:
    a. Align the pins of the support rails above and below the nut clip you installed in the previous step..
    b. On the right side of the rack, connect the flange on the end of the support rail to the front edge of the rack by using two M5 screws into the upper and lower screw holes, leaving the middle screw hole open. Ensure that the rail assembly is left slightly loose on the front of the rack to allow for the HMC to be inserted.
7. At the rear of the rack, on the right side, pull the free end of the support rail toward the rear and secure the flange of the support rail to the rack by using two M5 screws, leaving the middle screw hole open.
8. Repeat the step “6” on page 9 and step “7” on page 9 to install the right support rail assembly on the right side of the rack.
9. In front of the rack, install the HMC workstation into the rack by performing the following steps:
    a. Holding the HMC workstation level, insert the inner rails into the HMC support rails that you installed in the previous step. Push the HMC forward until the flanges on the front of the HMC are flush with the open screw holes on the front of the rack.

b. Connect the HMC to the left side of the frame using one M5 screw. Repeat this step on the right side of the rack.

Note: If present, remove the orange shipping brackets that are attached to the rear of the system and reinstall the screw back in.

10. Continue with “Installing the system into the rack and connecting and routing power cables” on page 10.

Installing the system into the rack and connecting and routing power cables

Install the system onto the rails and connect and route power cables.

About this task

⚠ **CAUTION:** This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

Procedure

1. Plug the power cords into the power supplies.

   Note: Do not connect the other end of the power cord to the power source now.

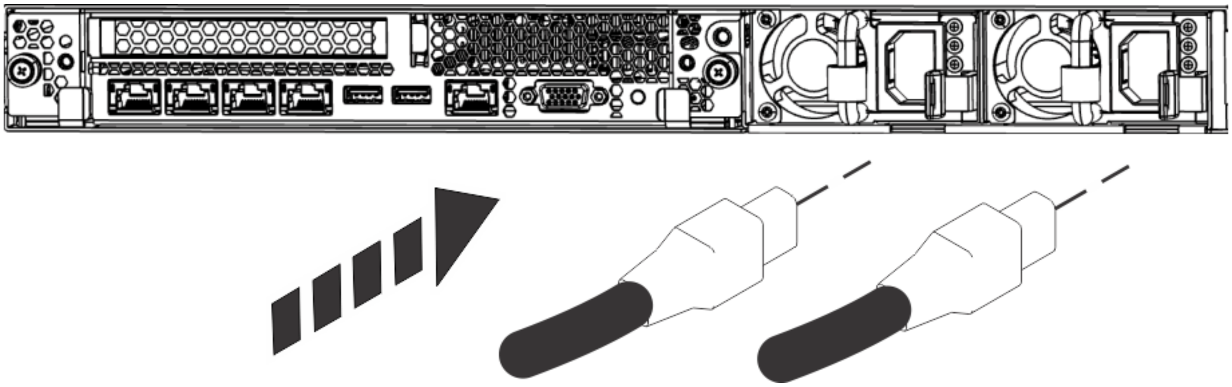

Figure 2. Plugging the power cords into the power supplies

2. Fasten the hook-and-loop fasteners to secure the power cords.
3. Continue with “Cabling the rack-mounted 7063-CR2 HMC” on page 10.

Cabling the rack-mounted 7063-CR2 HMC

Learn how to physically install your rack-mounted Hardware Management Console (HMC).

Procedure

1. Ensure that the HMC is installed into a rack and the power cords are plugged into the power supplies. For more information, see “Installing the system into the rack and connecting and routing power cables” on page 10. After you install the HMC into a rack, continue with the next step.
2. Connect the keyboard, monitor, and mouse.

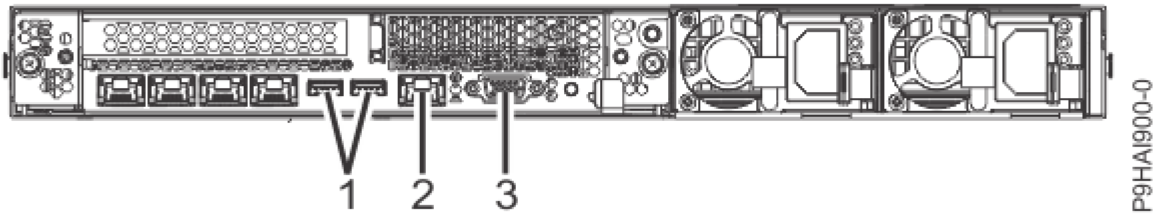

Figure 3. Rear ports

Table 4. Input and output ports

Identifier	Description
1	USB 2.0 used for keyboard and mouse
2	Ethernet Intelligent Platform Management Interface (IPMI)
3	Video Graphics Array (VGA) that is used for the monitor. Only the 1024 x 768 at 60 Hz VGA setting is supported. Only up to a 3-meter cable is supported.

Note: The system has two front USB ports that you can use.

3. Connect the Ethernet Intelligent Platform Management Interface (IPMI) port to a network.

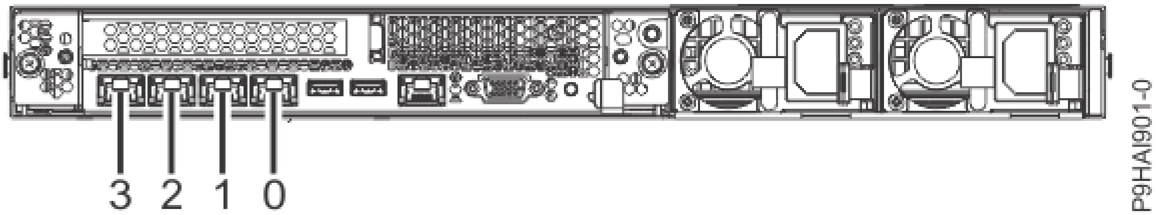

Figure 4. Ethernet ports

Table 5. Ethernet ports

Identifier	Description
0	Shared Ethernet Intelligent Platform Management Interface (IPMI) and HMC Network Connection
1, 2, and 3	HMC network connection

Notes:

- When the HMC is equipped with an optional 10 Gb PCI Express (PCIe) adapter, two additional ethernet ports are available.
- This connection is required to access the baseboard management controller (BMC) on the HMC. Access to the BMC is required for service tasks and to maintain the HMC firmware. For more information, see “Types of HMC network connections” on page 52.

Warning: This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Please contact IBM for more information.

4. Connect the Ethernet cable that is intended for the connection to the managed system or systems.

Notes:

- If you are using a shared connection for IPMI and HMC, a single cable to port 0 in Figure 2 can satisfy both requirements for IPMI and HMC.
- To learn more about the HMC network connections, see “HMC network connections” on page 51.

5. If your managed system is already installed, you can verify that the Ethernet cable connection is active by observing the green status lights at both the HMC and managed system Ethernet ports as your installation progresses.
6. Plug the system power cords and the power cords for any other attached devices into the alternating current (AC) power source.

7. Verify the power status by using the power supply LEDs as indicators. For more information, see LEDs on the 7063-CR2 systemLEDs on the 7063-CR2 system.
8. Press the power button to start the system. The power-on light stops flashing and remains on, indicating that the system power is on.

Results

Next, you need to install and configure your HMC software. Continue with “Configuring the 7063-CR2 HMC” on page 12.

Configuring the 7063-CR2 HMC

Learn how to install and configure the Hardware Management Console (HMC).

Check the HMC version that is shipped with your HMC. To find out how to view the HMC machine code version and release, see Check the HMC version that is shipped with your HMC. You can download the latest HMC version that is available from the Fix Central website. Use removable media (such as a DVD or USB) to create a bootable ISO file from the HMC package (ISO image).

Note: The following table describes the predefined (default) login information for the HMC and BMC interfaces.

Table 6.

Console or Interface	Default ID	Default Password	Description
BMC (OpenBMC)	`root`	`0penBmc`	The root user ID and password are used to log in to the BMC for the first time.
HMC	`hscroot`	`abc123`	The hscroot user ID and password are used to log in to the HMC for the first time. They are case-sensitive and can only be used by a member of the super administrator role.
HMC	`root`	`passw0rd`	The root user ID and password are used by the service provider to perform maintenance procedures. They cannot be used to log in to the HMC.

Note: The following installations are shown as examples.

Installing the HMC by using USB flash drive

To install the HMC by using USB flash drive, complete the following steps for Linux® systems:

Note: For examples in different operating systems, see:

- Windows: USB flash installation media (Windows)
- Mac: USB flash installation media (macOS)

1. Download the HMC version that you want from the Fix Central website.
2. Run the following command to identify the device name of the USB drive when it is plugged in: **`lsblk`**.

For example: **`/dev/sdb`** (where **sdb** is the name of the USB drive)

3. Run the following command to wipe the USB drive: **`wipefs --all /dev/sdX`**.

   For example: **`wipefs --all /dev/sdb`**

4. Run the following command to verify the size of the disk under the SIZE column: **`lsblk`**.

   For example: When a 16 GB USB drive shows as 14.3 GB, round it down to 14 GB for the next step "5" on page 13.

5. Run the following command to format the disk and create a partition: **`parted /dev/sdX`**

   From the parted utility, run the following three commands:

   ```
   mklabel gpt
   mkpart primary ext3 1MiB <size>GiB
   quit
   ```

   Note: size is the size of the USB drive obtained in the step "4" on page 13.

   For example:

   ```
   parted /dev/sdb
   mklabel gpt
   mkpart primary ext3 1MiB 14GiB
   quit
   ```

6. Run the following command to copy the ISO onto the partition: **`cat HMC-Recovery-ppc64le.iso > /dev/sdX1`**.

   For example: **`cat HMC-9.2.950.0-2103300827-ppc64le.iso > /dev/sdb1`**

7. Insert the USB drive, and power on the system.

   Note: The USB drive must be at least 8 GB. Certain USB drives might be too wide to fit properly into the USB port at the rear of the system. Test the fit of your USB drive before you proceed.

8. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **USB**.

Installing the HMC by using virtual media from the BMC

To install the HMC by using virtual media from the BMC, complete the following steps:

1. Open a supported web browser. In the address bar, enter the IP address of the BMC that you want to connect to. For example, you can use the format `https://<BMC IP>` in the address bar of the web browser.
2. From the **OpenBMC logon** window, enter the **Host** address of the BMC and the **Username** and **Password** that is assigned to you.

   Note: The default user ID is `root` and the default password is `0penBmc`.

   If you are using firmware level OP940.01, or later, the root password is expired by default. You must change the default password before you can access the BMC. For more information about changing the expired default password, see Setting the password.

   If you forgot your password, you can perform a factory reset of the system to restore the default password. To reset the system, see Performing a factory reset.
3. Click **Log in**.
4. Select **Server control**.
5. Select **Virtual Media**.
6. Click **Choose file**.
7. Locate the HMC Recovery media ISO and click **Open**.

8. Click **Start**.
9. Power on the system.
10. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **USB**.

Installing the HMC by using an external USB attached DVD drive

To install the HMC by using an external USB attached DVD drive, complete the following steps:

1. Download the HMC recovery version that you want from the Fix Central website.
2. Burn the HMC recovery DVD image to a DVD-R DL media as an image.
3. Power off the HMC.
4. Connect the external USB DVD drive to the HMC and insert the HMC recovery DVD.

   Note: You might need to connect the USB DVD drive to an external power source or use a USB Y cable to connect to an extra USB port to provide sufficient power to the DVD drive.
5. Power on the HMC.

   Note: The display monitor might show no signal during startup. The process might take 2 or 3 minutes before the display monitor shows any status.
6. When the Petitboot bootloader starts, navigate to stop the automatic boot.

   Note: A 10-second timeout is enforced. If no action is taken within 10 seconds, the system attempts to boot from the hard disk drive.
7. Wait until the **CD/DVD** device appears in the Petitboot menu.

   Note: This process can take up to a minute.
8. Select the **Install Hardware Management Console** option that is located under **CD/DVD**.

Installing the 7063-CR1 into a rack

Learn how to install the 7063-CR1 Hardware Management Console (HMC) into a rack.

You can view the online installation documentation, or you can print the PDF version of the same information. To view or print the PDF version, see Installing and configuring the Hardware Management Console.

Prerequisites for installing the rack-mounted 7063-CR1 system

Use the information to understand the prerequisites that are required for installing the system.

About this task

⚠

CAUTION: 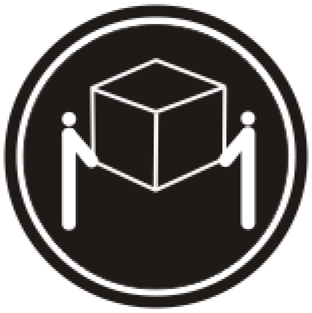 or 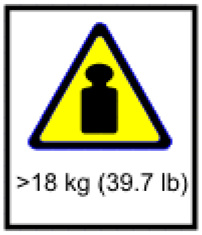

or 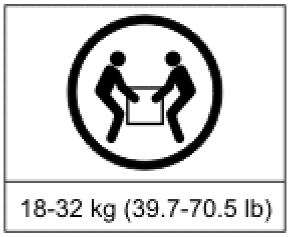

The weight of this part or unit is between 18 and 32 kg (39.7 and 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

You might need to read the following documents before you begin to install the server:

- The latest version of this document is maintained online, see Installing the 7063-CR1 into a rack (http://www.ibm.com/support/knowledgecenter/POWER9/p9hai/p9hai_install7063_kickoff.htm).
- To plan your server installation, see Site and hardware planning.

Procedure

Ensure that you have the following items before starting your installation:

- Size 2 Phillips screwdriver
- Flat-head screwdriver
- Box cutter
- Electrostatic discharge (ESD) wrist strap
- Rack with one Electronic Industries Association (EIA) unit (1U) of space.

Note: If you do not have a rack that is installed, install the rack. For instructions, see Racks and rack features (http://www.ibm.com/support/knowledgecenter/POWER9/p9hbf/p9hbf_9xx_kickoff.htm).

Completing inventory for your system

Use this information to complete inventory for your system.

Procedure

1. Verify that you received all the boxes you ordered.
2. Unpack the server components as needed.
3. Complete a parts inventory and verify that you have received all the parts that you ordered before you install each server component.

   Note:

   Your order information is included with your product. You can also obtain order information from your marketing representative or the IBM Business Partner.

   If you have incorrect, missing, or damaged parts, consult any of the following resources:

   - Your IBM reseller.
   - IBM Rochester manufacturing automated information line at 1-800-300-8751 (United States only).
   - The Directory of worldwide contacts website (http://www.ibm.com/planetwide). Select your location to view the service and support contact information.

Determining and marking the location in the rack for the 7063-CR1 system

You might need to determine where to install the system unit into the rack.

Procedure

1. Read the Rack safety notices (http://www.ibm.com/support/knowledgecenter/POWER9/p9hbf/p9hbf_racksafety.htm).
2. Determine where to place the system unit in the rack. As you plan for installing the system unit in a rack, consider the following information:

   - Organize larger and heavier units into the lower part of the rack.
   - Plan to install system units into the lower part of the rack first.
   - Record the Electronic Industries Alliance (EIA) locations in your plan.
3. If necessary, remove the filler panels to allow access to the inside of the rack enclosure where you plan to place the unit, as shown in Figure 5 on page 16.

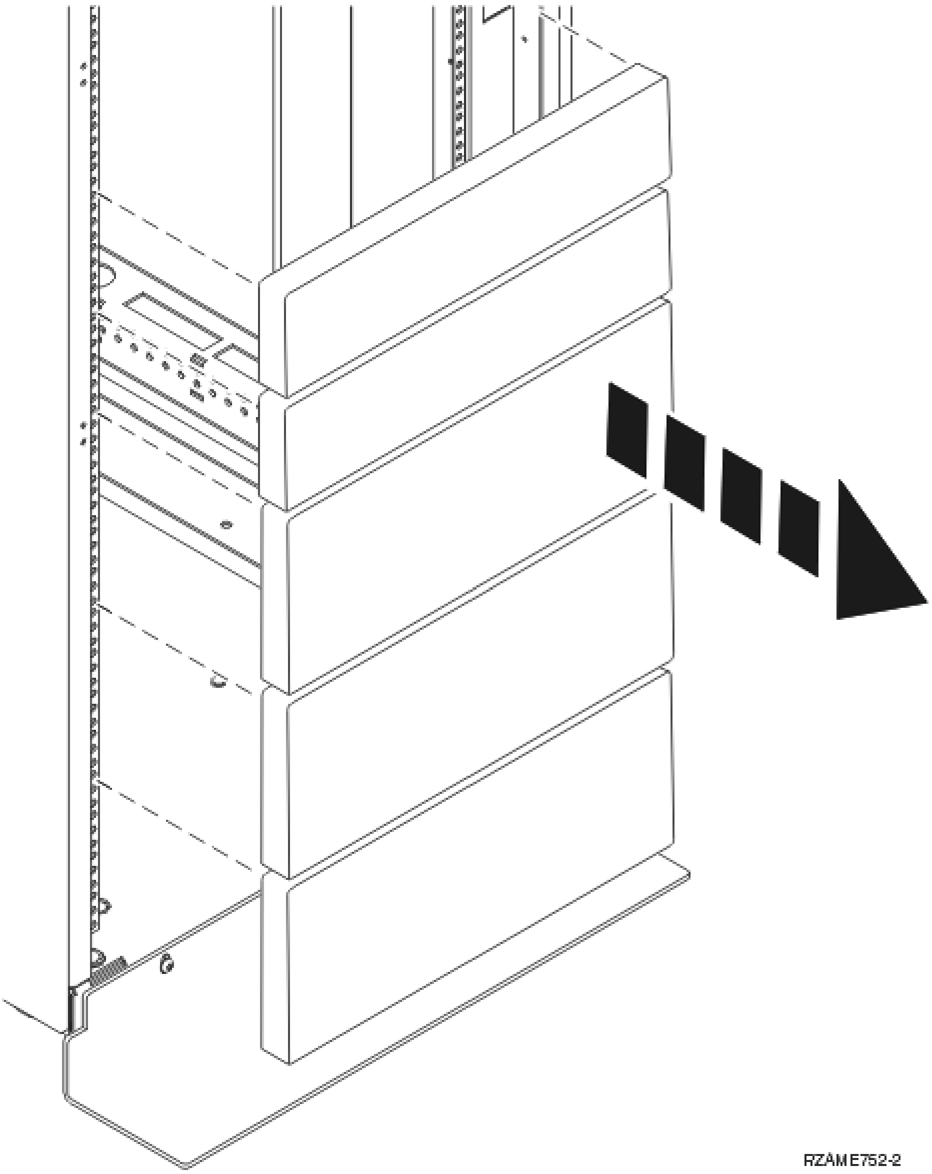

Figure 5. Removing the filler panels

4. Determine to place the system in the rack. Record the EIA location.
5. Facing the front of the rack and working from the right side, use tape, a marker, or pencil to mark the lower hole of each EIA unit.
6. Repeat step “5” on page 16 for the corresponding holes located on the left side of the rack.
7. Go to the rear of the rack.
8. On the right side, find the EIA unit that corresponds to the bottom EIA unit marked on the front of the rack.
9. Mark the bottom EIA unit.
10. Mark the corresponding holes on the left side of the rack.

Attaching the fixed rails to the system chassis and to the rack

You must install the rails onto the chassis and into the rack. Use this procedure to perform this task.

About this task

⚠ **Attention:** To avoid rail failure and potential danger to yourself and to the unit, ensure that you have the correct rails and fittings for your rack. If your rack has square support flange holes or screw-thread support flange holes, ensure that the rails and fittings match the support flange holes that are used on your rack. Do not install mismatched hardware by using washers or spacers. If you do not have the correct rails and fittings for your rack, contact your IBM reseller.

Note: The system requires 1 EIA rack unit (1U) of space.

Ensure that you have the necessary parts to install the rails. The following parts are included with the rail kit:

- Slide rail screws, used to attach the two parts of each slide rail
- Slide rail rack screws, used to secure the rails to the rack
- Rails
- 10 - 32 x 0.635 cm (0.25 in.) screws, used to attach the rails to system chassis

Procedure

1. Remove the rail pieces from the packaging and put them on a work surface.
2. Replace the rail rack square pins (**A**) and (**D**) with the rail rack round pins.
3. Connect the two parts of each rack slide rail. To connect the two parts of the rack slide rail, perform the following tasks:
   a. Identify the two pieces of the left rack slide rail. Align the short and long pieces (**C**). Ensure that the rack rail pins are pointing in the same direction (**A**) and (**D**).

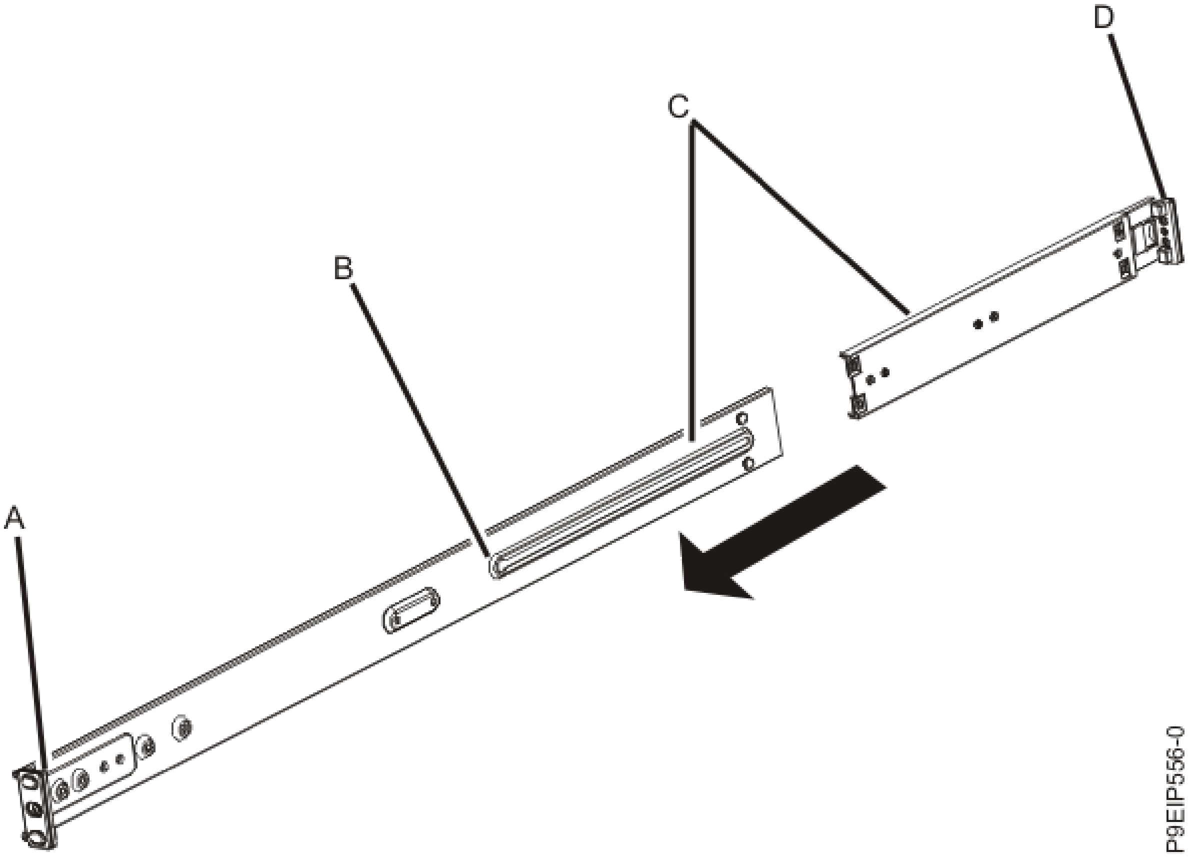

   b. The shorter piece of the rack slide rail has a metal pin. Insert the pin into the hole in the longer piece of the rack slide rail (**B**). Slide the shorter piece of the rack rail into the longer piece of the rack rail.
   c. Align the holes in the two pieces of the rack slide rails. Using a Philips-head screwdriver, attach the two parts by loosely screwing two threaded rail screws through the holes in the rack slide rail.

   Note: Do not tighten the rack slide rail screws.

   d. Repeat these steps for the right slide rail.

4. Install the rack slide rails into the rack.

   a. Move to the front of the rack.

   b. Select the left rack slide rail, and locate the EIA unit that you previously marked. Each slide rail is also marked **Back**, to designate the rear of the rack. Ensure that you are holding the front end of the rack slide rail.

   c. Extend the rail from the front of the rack to the back of the rack and align the rack slide rail pins with the holes in the rack flange that you previously marked.

   d. Push the rack rail pins into the rear rack flange until the rear rack rail latch clicks into place.

   e. Pull the front of the rack rail toward the front of the rack rail flange. Align the slide rail pins with the holes in the rail flange and pull them until the rail latch clicks into place.

   f. Using a screwdriver, tighten the rail screws that you installed in step 2.

   Note: You might need 2U of space to access and tighten the rail screws.

   g. Repeat steps 4a - 4f for the right slide rail.

Installing the system into the rack and connecting and routing power cables

Install the system onto the rails and connect and route power cables.

About this task

CAUTION: 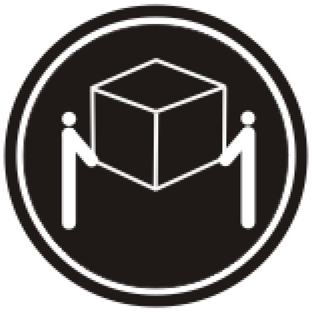 or 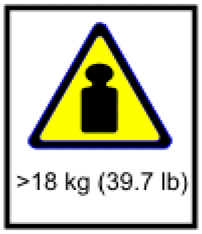

or 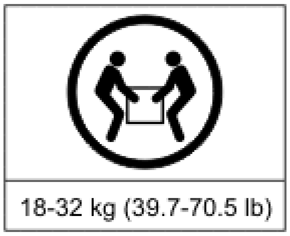

The weight of this part or unit is between 18 and 32 kg (39.7 and 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

Procedure

1. Remove the protective plastic film from the top of the system chassis.
2. Move to the front of the rack.
3. Using two people, one on each side of the system, lift the system and align the system chassis rails on each side of the chassis with the rack slide rails.
4. Gently push the system toward the rear of the rack.
5. Secure the system to the rack by screwing a screw with washer through the handles on each side of the system chassis.

   Note: You must use washers with the screws. Slide a washer onto to each of the two longer screws (1.5 cm (0.59 in.)) that is included with the rail kit. Screw the screw with the washer through the right and left side of the system in the front.

6. Plug the power cords into the power supplies.

   Note: Do not connect the other end of the power cord to the power source now.

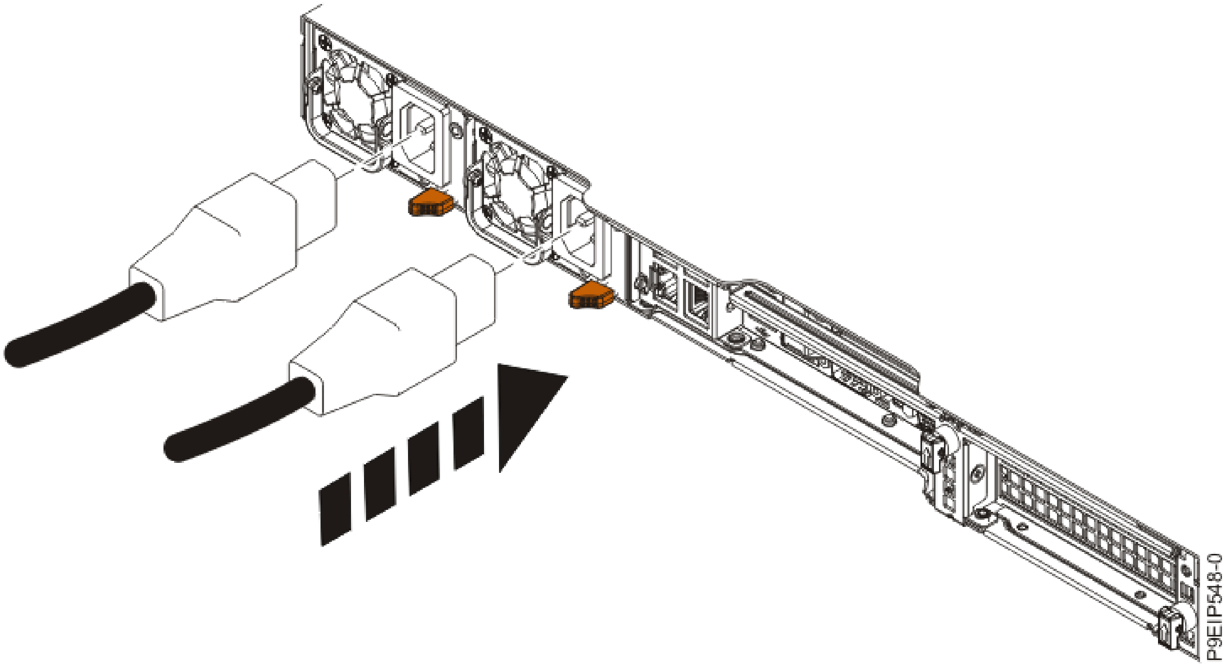

Figure 6. Plugging the power cords into the power supplies

7. Continue with "Cabling the rack-mounted 7063-CR1 HMC" on page 19.

Cabling the rack-mounted 7063-CR1 HMC

Learn how to physically install your rack-mounted Hardware Management Console (HMC).

Procedure

1. Ensure that the HMC is installed into a rack and the power cords are plugged into the power supplies. For more information, see "Installing the system into the rack and connecting and routing power cables" on page 18. After you install the HMC into a rack, continue with the next step.

   Note: If a plug is covering a port that you need to use on the rear of the system, remove and discard it. The port covers ensure that you are reminded that you must reset the Administrator password on your managed system upon initial system IPL.

2. Connect the keyboard, monitor, and mouse.

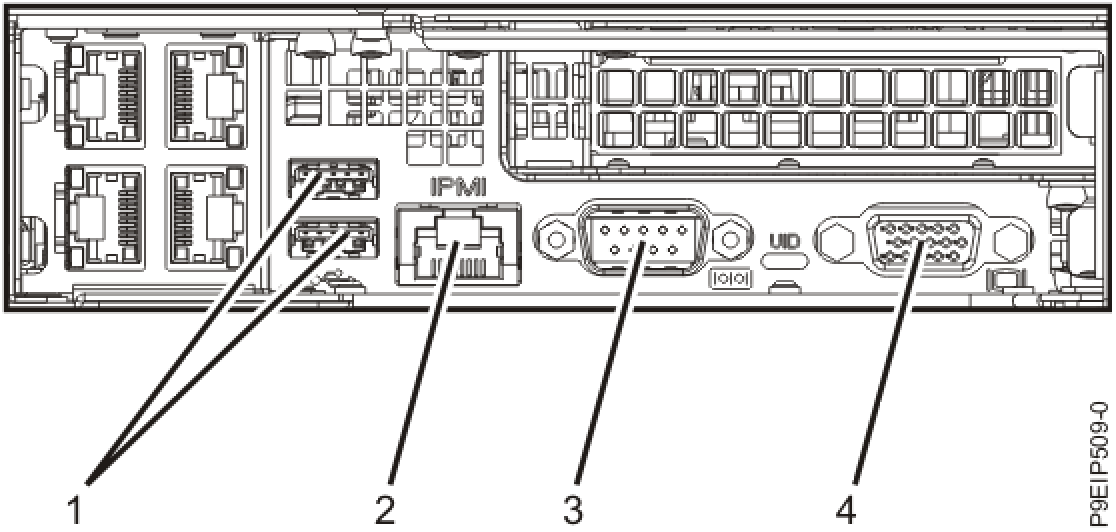

Figure 7. Rear ports

Table 7. Input and output ports

Identifier	Description
1	USB 2.0 used for keyboard and mouse
2	Ethernet Intelligent Platform Management Interface (IPMI)
3	Serial IPMI
4	Video Graphics Array (VGA) that is used for the monitor. Only the 1024 x 768 at 60 Hz VGA setting is supported. Only up to a 3-meter cable is supported.

Note: The system has two front USB ports that you can use. The front serial port is non-functional.

3. Connect the Ethernet cable that is intended for the connection to the managed system or systems.

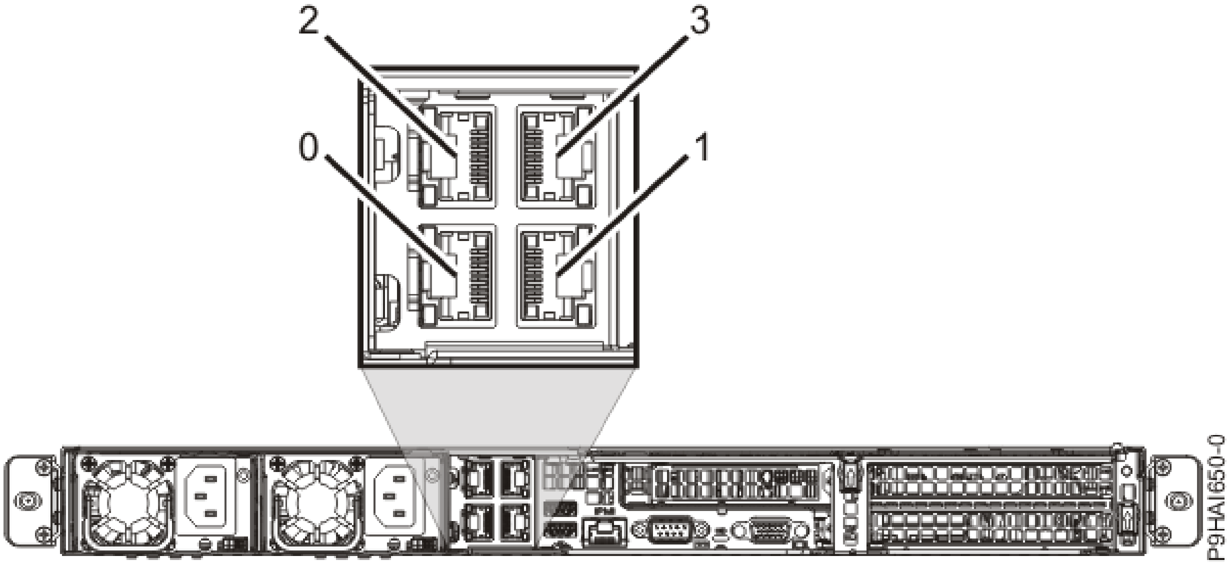

Figure 8. Ethernet ports

Note: To learn more about the HMC network connections, see "HMC network connections" on page 51.

4. If your managed system is already installed, you can verify that the Ethernet cable connection is active by observing the green status lights at both the HMC and managed system Ethernet ports as your installation progresses.
5. Connect the Ethernet Intelligent Platform Management Interface (IPMI) port to a network.

   Note: This connection is required to access the baseboard management controller (BMC) on the HMC. Access to the BMC is required for service tasks and to maintain the HMC firmware. For more information, see "Types of HMC network connections" on page 52.

6. Plug the system power cords and the power cords for any other attached devices into the alternating current (AC) power source.
7. Verify the power status by using the power supply LEDs as indicators. For more information, see LEDs on the 7063-CR1 systemLEDs on the 7063-CR1 system.

Results

Next, you need to install and configure your HMC software. Continue with "Configuring the 7063-CR1 HMC" on page 21.

Configuring the 7063-CR1 HMC

Learn how to install and configure the Hardware Management Console (HMC).

Check the HMC version that is shipped with your HMC. You can download the latest HMC version that is available from the Fix Central website. Use removable media (such as a DVD or USB) to create a bootable ISO file from the HMC package (ISO image).

Note: The following table describes the predefined (default) login information for the HMC and BMC interfaces.

Table 8.

Console or Interface	Default ID	Default Password	Description
BMC	ADMIN	ADMIN	The ADMIN user ID and password are used to log in to the BMC for the first time.
HMC	hscroot	abc123	The hscroot user ID and password are used to log in to the HMC for the first time. They are case-sensitive and can only be used by a member of the super administrator role.
HMC	root	passw0rd	The root user ID and password are used by the service provider to perform maintenance procedures. They cannot be used to log in to the HMC.

Note: The following installations are shown as examples.

Installing the HMC by using USB flash drive

To install the HMC by using USB flash drive, complete the following steps for Linux systems:

Note: For examples in different operating systems, see:

- Windows: USB flash installation media (Windows)
- Mac: USB flash installation media (macOS)

1. Download the HMC version that you want from the Fix Central website.
2. Run the following command: **`dd bs=4M if=/path/to/HMC_ISO_FILE.iso of=/dev/sdx status=progress && sync`** (where **sdx** is the name of the USB drive).

   Note: You can run the Linux command `lsblk` to determine the device name of the USB drive when it is plugged in.
3. Insert the USB drive, and power on the system.

   Note: The USB drive must be at least 4 GB. Certain USB drives might be too wide to fit properly into the USB port at the rear of the system. Test the fit of your USB drive before you proceed.
4. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **USB**.

Installing the HMC by using remote media from the console viewer

To install the HMC by using remote media from the console viewer, complete the following steps:

1. Log in to the BMC web interface (`http://<bmc-ip>`).
2. Select **Remote Control**.
3. Select **Console Redirection**.
4. Click **Launch Console**.
5. In the Java™ iKVM Viewer, select **Virtual Media** > **Virtual Storage**.
6. Under **Logical Drive Type**, select **ISO File**.
7. Click **Open Image** and locate the ISO file on your system.
8. Press **Plugin** to mount the ISO file.
9. Power on the system.
10. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **CD/DVD**.

Installing the HMC by using an external USB attached DVD drive

To install the HMC by using an external USB attached DVD drive, complete the following steps:

1. Download the HMC recovery version that you want from the Fix Central website.
2. Burn the HMC recovery DVD image to a DVD-R media as an image. Alternatively, you can order the recovery media on DVD.
3. Power off the HMC.
4. Connect the external USB DVD drive to the HMC and insert the HMC recovery DVD.

   Note: You might need to connect the USB DVD drive to an external power source or use a USB Y cable to connect to an extra USB port to provide sufficient power to the DVD drive.
5. Power on the HMC.

   Note: The display monitor might show no signal during startup. The process might take 2 or 3 minutes before the display monitor shows any status.
6. When the Petitboot bootloader starts, navigate to stop the automatic boot.

   Note: A 10-second timeout is enforced. If no action is taken within 10 seconds, the system attempts to boot from the hard disk drive.
7. Wait until the **CD/DVD** device appears in the Petitboot menu.

   Note: This process can take up to a minute.
8. Select the **Install Hardware Management Console** option that is located under **CD/DVD**.

Installing the HMC by using remote media that is hosted by an SMB file server

To install the HMC by using remote media that is hosted by a Server Message Block (SMB) file server, complete the following steps:

1. Copy the recovery ISO file to a share host on your SMB-compliant file server.

   Note: Server Message Block version 3 (SMBv3) is not supported.
2. Log in to the BMC web interface (`http://<bmc-ip>`).
3. Select **Virtual Media**.
4. Select **CD-ROM Image**.
5. Complete the following information:

Share host

The IP of the SMB host. If you are using the host name, ensure that the domain name system (DNS) on the BMC is correctly configured.

Path to image

The SMB path to the system. For example: `/<share name>/<rest of path>/<name of iso>.iso`

User (optional)

The user name that is used to log in to the SMB host.

Password (optional)

The password for the user.

6. Click **Save**.
7. Click **Mount**.
8. Device 1 now shows the following message: **There is an iso file mounted.**

   Note: If the message does not appear, recheck the information and repeat steps 6 - 8.
9. Power on the system.
10. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **CD/DVD**.

Optional: Update the HMC firmware level by using the included USB memory key

Note: If your configuration included an HMC firmware update on a USB memory key, complete the following steps to update the HMC firmware level.

To update the HMC firmware level by using the included USB memory key, complete the following steps:

1. Insert the USB memory key drive into the USB port at the rear of the system.
2. Power on the system and log on to the HMC.
3. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.
4. In the content pane, click **Update the Hardware Management Console**.
5. Follow the onscreen instructions in the Install HMC Corrective Service wizard.

Next, you need to configure your HMC software. For instructions, see “Configuring the HMC” on page 51.

Related concepts

Configure BMC connectivity (7063-CR1)
You can configure or view the network settings on the BMC for the management console.

Installing the 7042-CR9 HMC into a rack

Learn how to install the 7042-CR9 Hardware Management Console (HMC) into a rack.

Before you begin

Complete a parts inventory. The following illustrations show the items that you need to install the server in the rack cabinet. If any items are missing or damaged, contact your place of purchase.

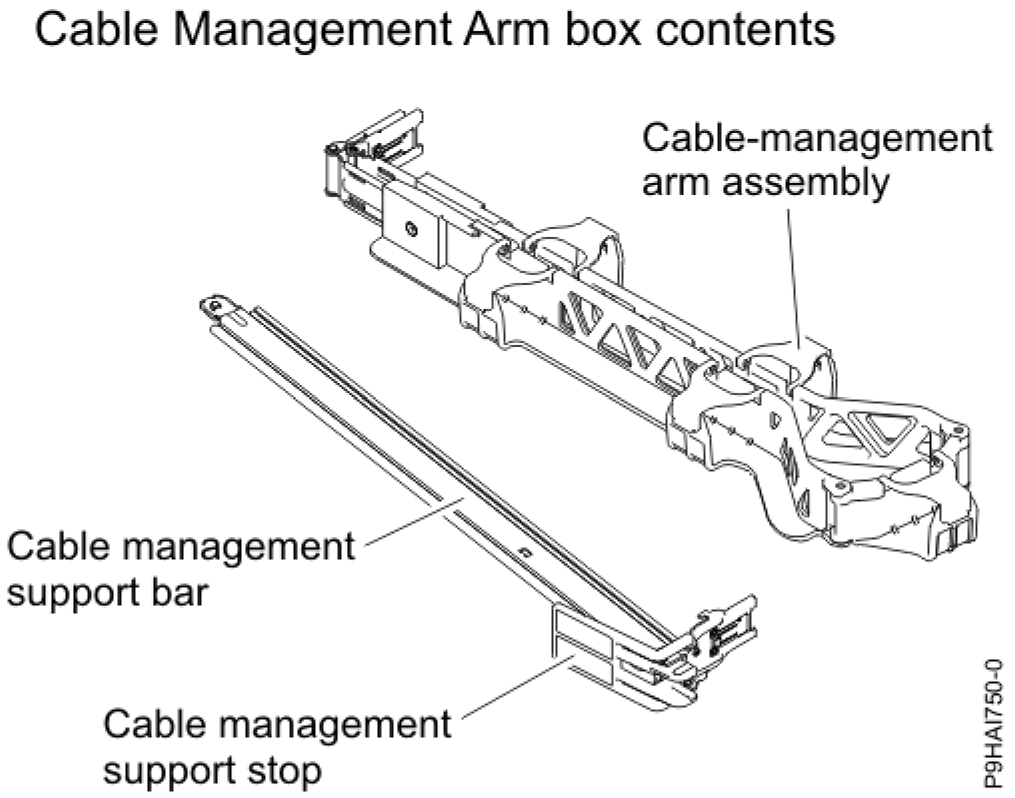
Cable Management Arm box contents
Cable-management arm assembly
Cable management support bar
Cable management support stop
P9HAI750-0

Figure 9. Cable management arm box contents

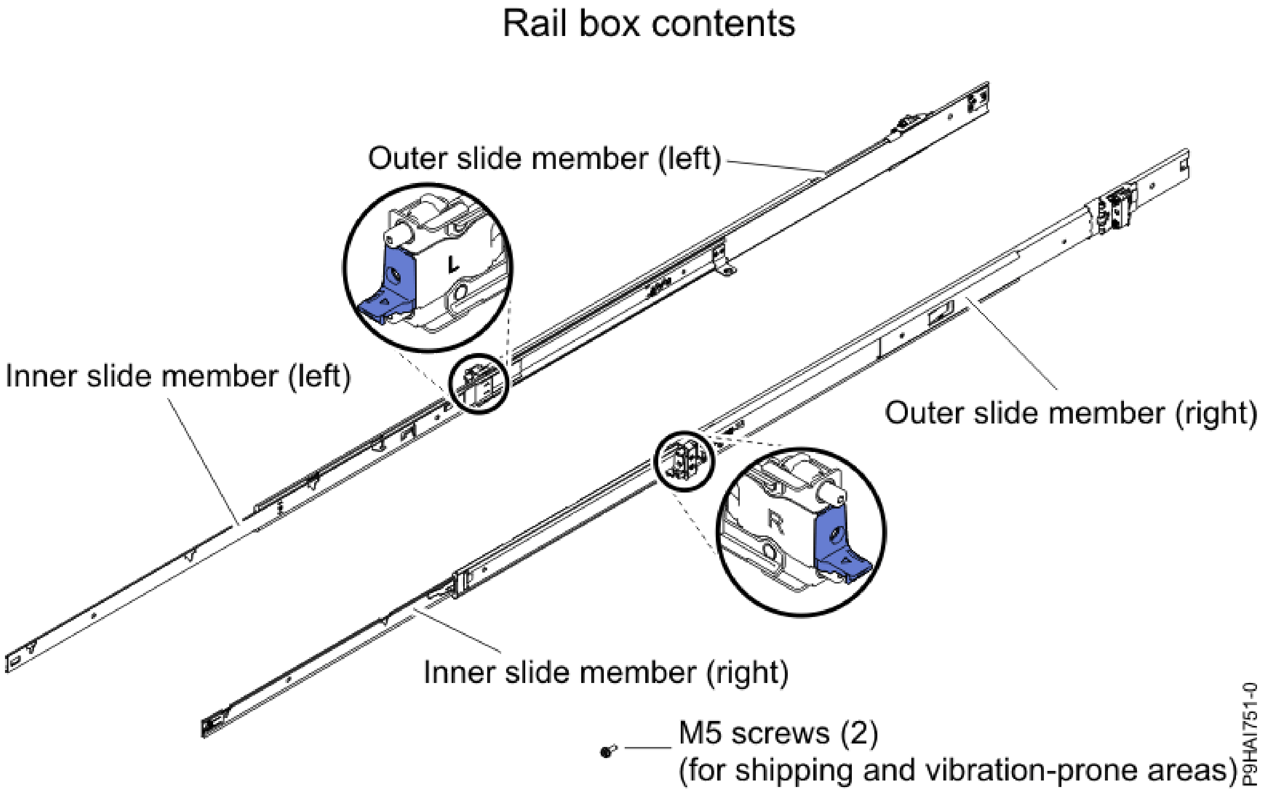
Rail box contents
Outer slide member (left)
L
Inner slide member (left)
Outer slide member (right)
R
Inner slide member (right)
M5 screws (2)
(for shipping and vibration-prone areas)
P9HAI751-0

Figure 10. Rail box contents

Note: You need both the slide rail box and the cable management arm box for this installation.

About this task

To install a 7042-CR9 HMC into a rack, complete the following steps:

Procedure

1. Select an available space (depending on the server you are installing) in your rack to install your server.

   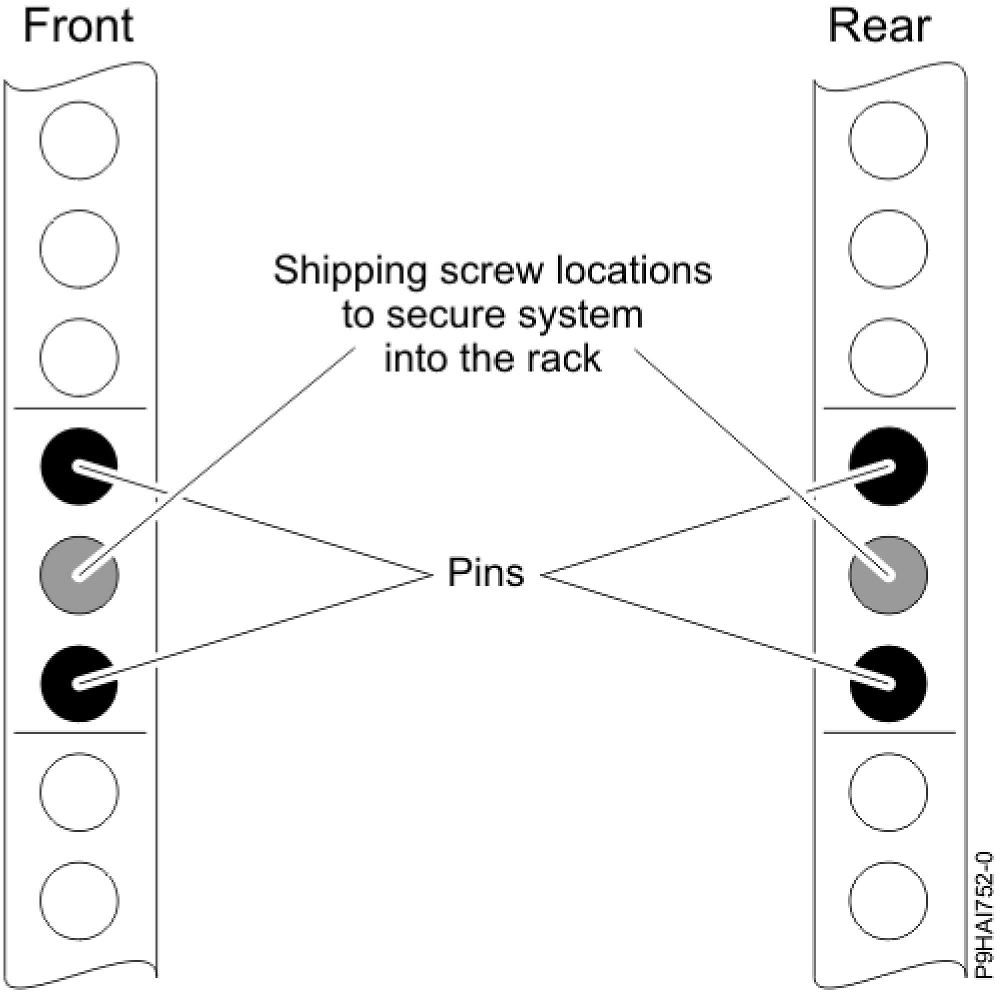

   Figure 11. Identifying a rack space

   Note: You need one unit (1 U) of space and the slide rails are installed in the bottom unit (U) of the one unit of space.
2. Extend the outer slide member all the way back until you hear an audible click. The rear rack mount bracket is now rotated into the unlocked position.

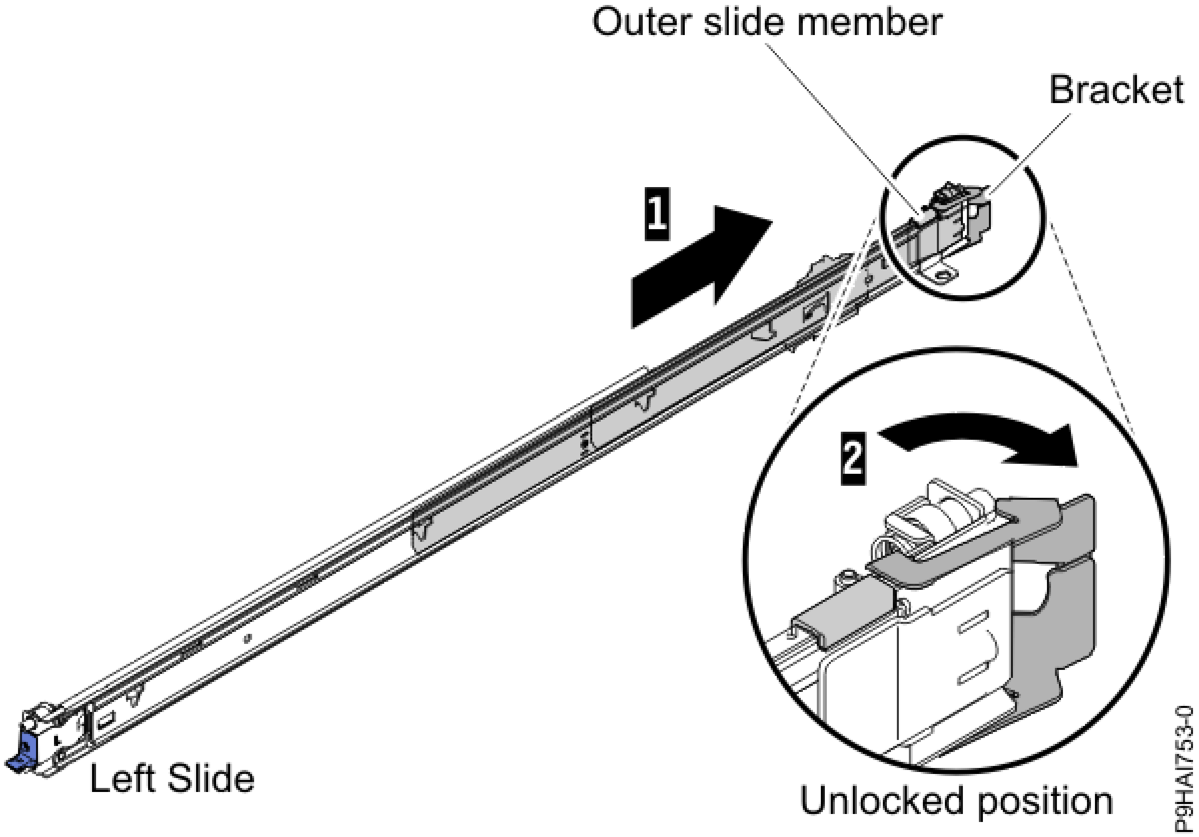

Figure 12. Slide rail and the outer slide member

Note: Each slide rail is marked as **R (right)** or **L (left)** on its end.

3. Align the rear end of the outer slide member against the holes on the rear of the rack. Line up the pins and push the slide in so that the pins go into the holes. The two slide pins protrude through the top and bottom holes on the EIA flange. Push the slide towards the rear of the rack until the rear rack mount bracket locks into place.

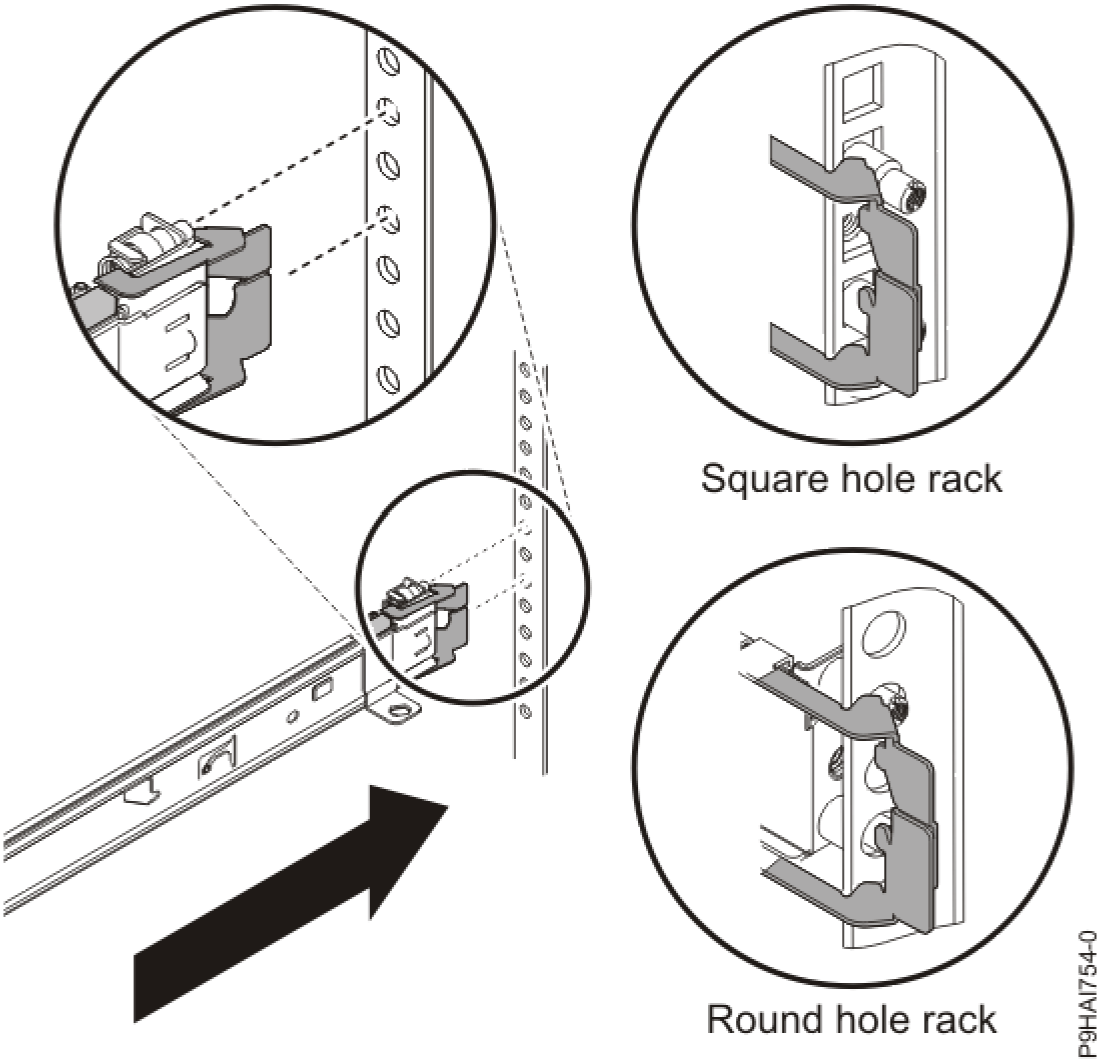

Figure 13. Align the pins with the holes in the rear of the rack

4. Rotate the front latch to the open position and align the front end of the outer slide member against the holes on the front of the rack. Line up the pins with holes in the EIA flanges and pull the slide forward so that the pins protrude through the holes. Lock the front of the slide by allowing the front latch to rotate to the closed position. Repeat steps 2- 4 for the other outer slide member.

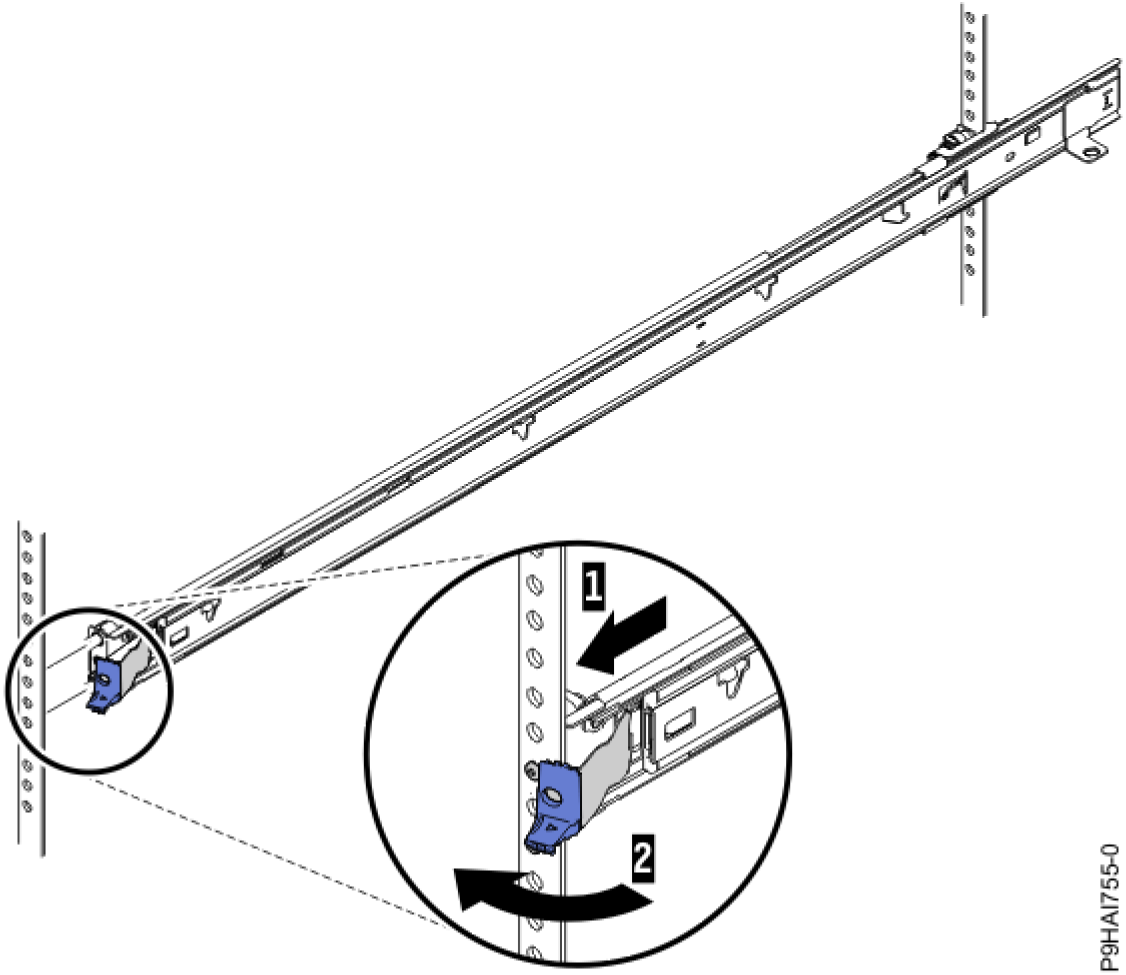

Figure 14. Front slide rail latch

5. Press on the release latches **(1)**. When you move the rack cabinet, or if you install the rack cabinet in a vibration-prone area, tighten the captive M5 screws **(2)** in the front of the server.

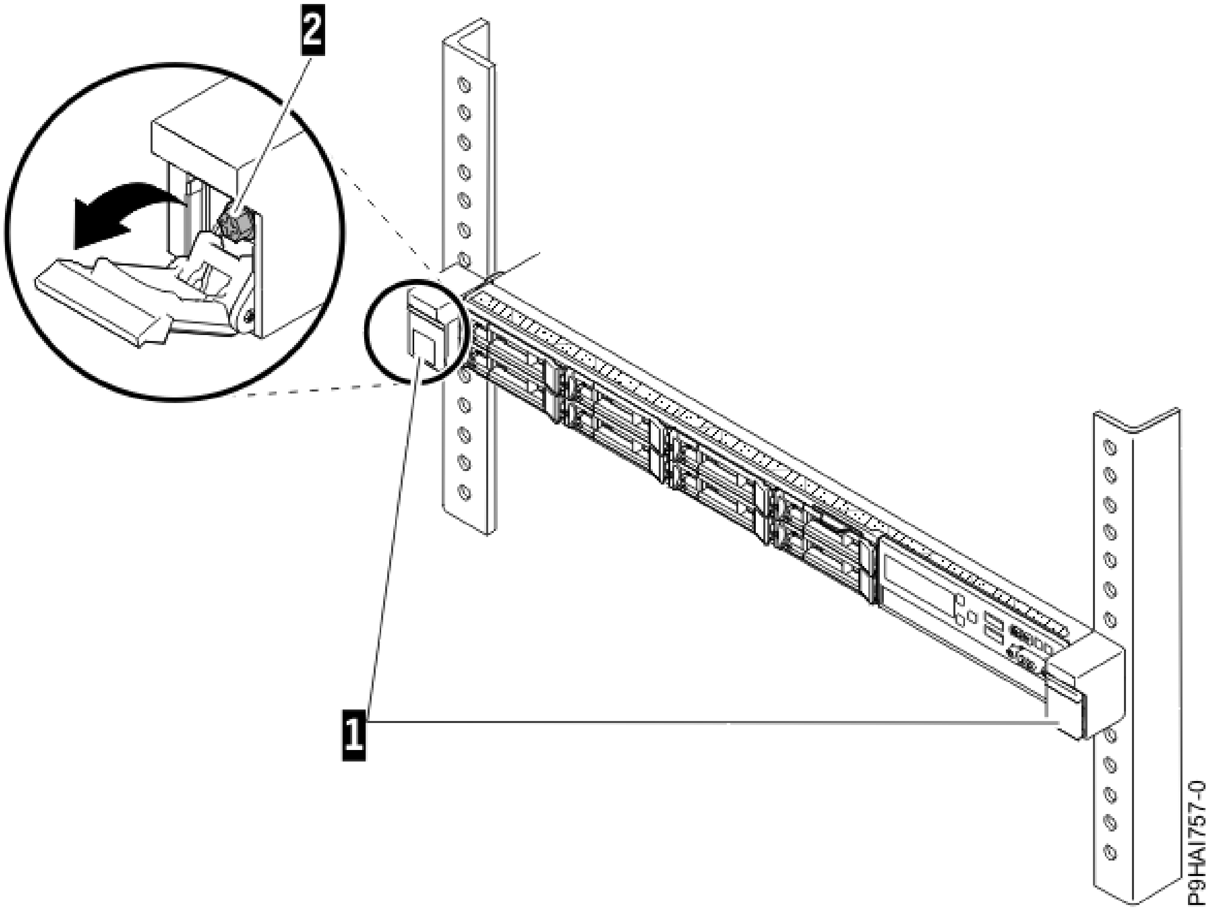

Figure 15. Rack front rail and pins

6. Pull the slide rails forward **(1)** until they click, twice, into place. Carefully lift the server and tilt it into position over the slide rails so that the rear nail heads **(2)** on the server line up with the slots in the slide rails. Lower the server down until the rear nail heads slide into the two rear slots, and then slowly lower the front of the server **(3)** until the other nail heads go into the other slots on the slide rails. Ensure that the front latch covers the front nail head so that the system is secured to the slide rails.

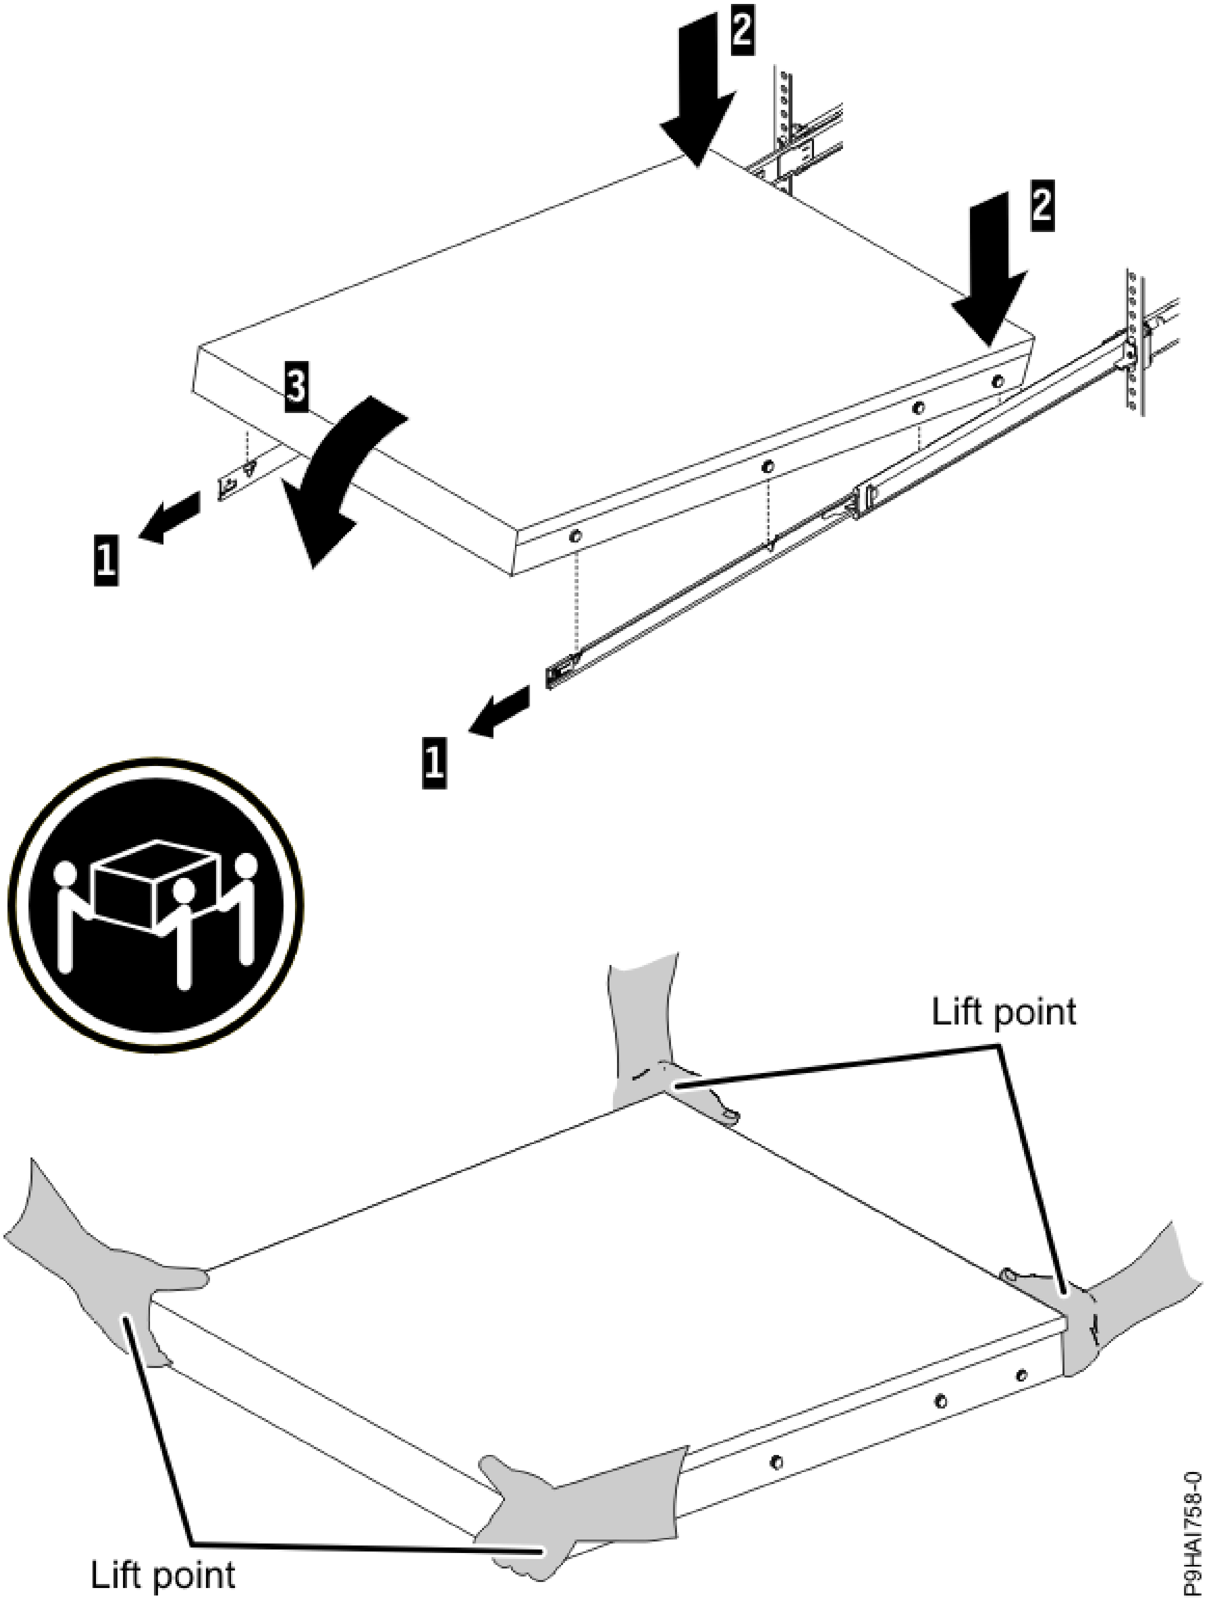

Figure 16. Slide rails extended, server nail heads aligned with slots in rail, and lift points

Note: Use safe practices while lifting. If you are installing a 1 U server, ensure that you have two people when you lift the server. Their hands must be positioned as illustrated in Figure 16 on page 30.

7. Lift the locking levers **(1)** on the slide rails and push the server **(2)** all the way into the rack until it clicks into place.

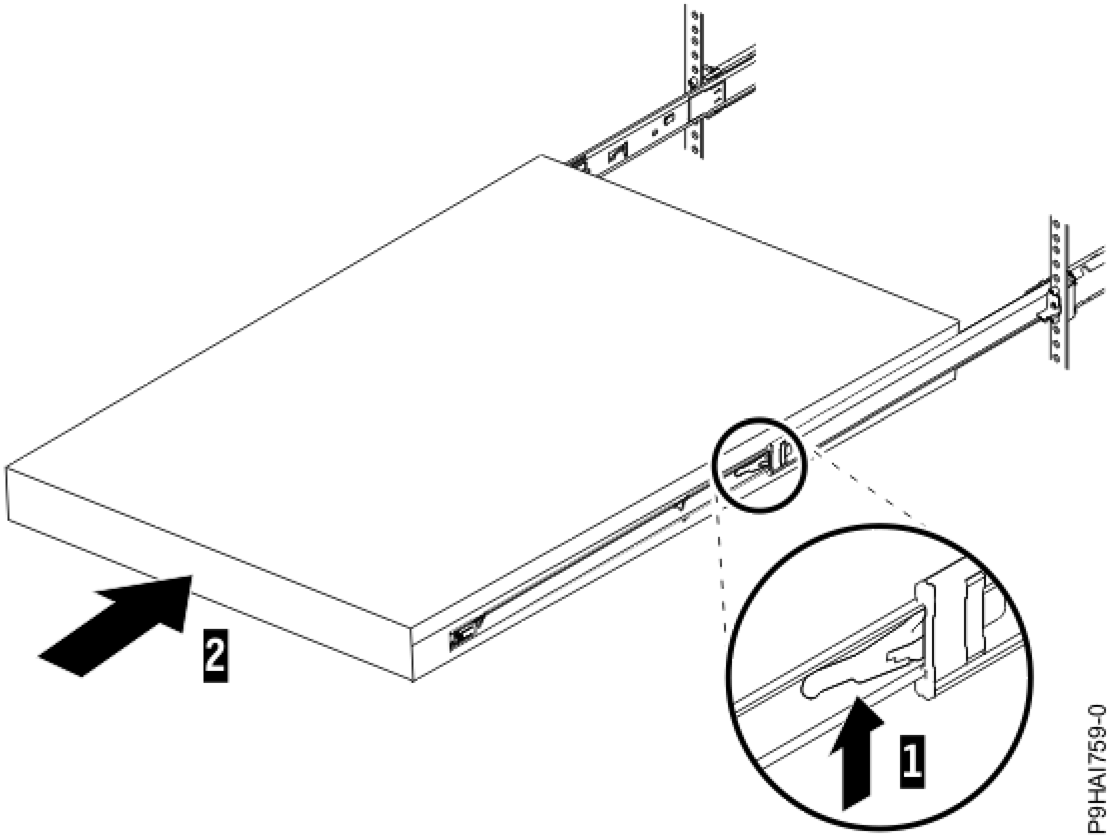

Figure 17. Release latches and server

8. The cable-management arm can be installed on either side of the server. Figure 18 on page 32 shows it being installed on the left side. It is best to install the cable-management arm so that it hinges on the side opposite to the power supplies to provide access to the power supplies. To install the cable-management arm on the right side, follow the instructions and install the hardware on the opposite side. Place the pin down **(1)** into the horizontal slot on the rear of the slide rails. Then rotate the other end of the bar toward the rack **(2)** toward the rack.

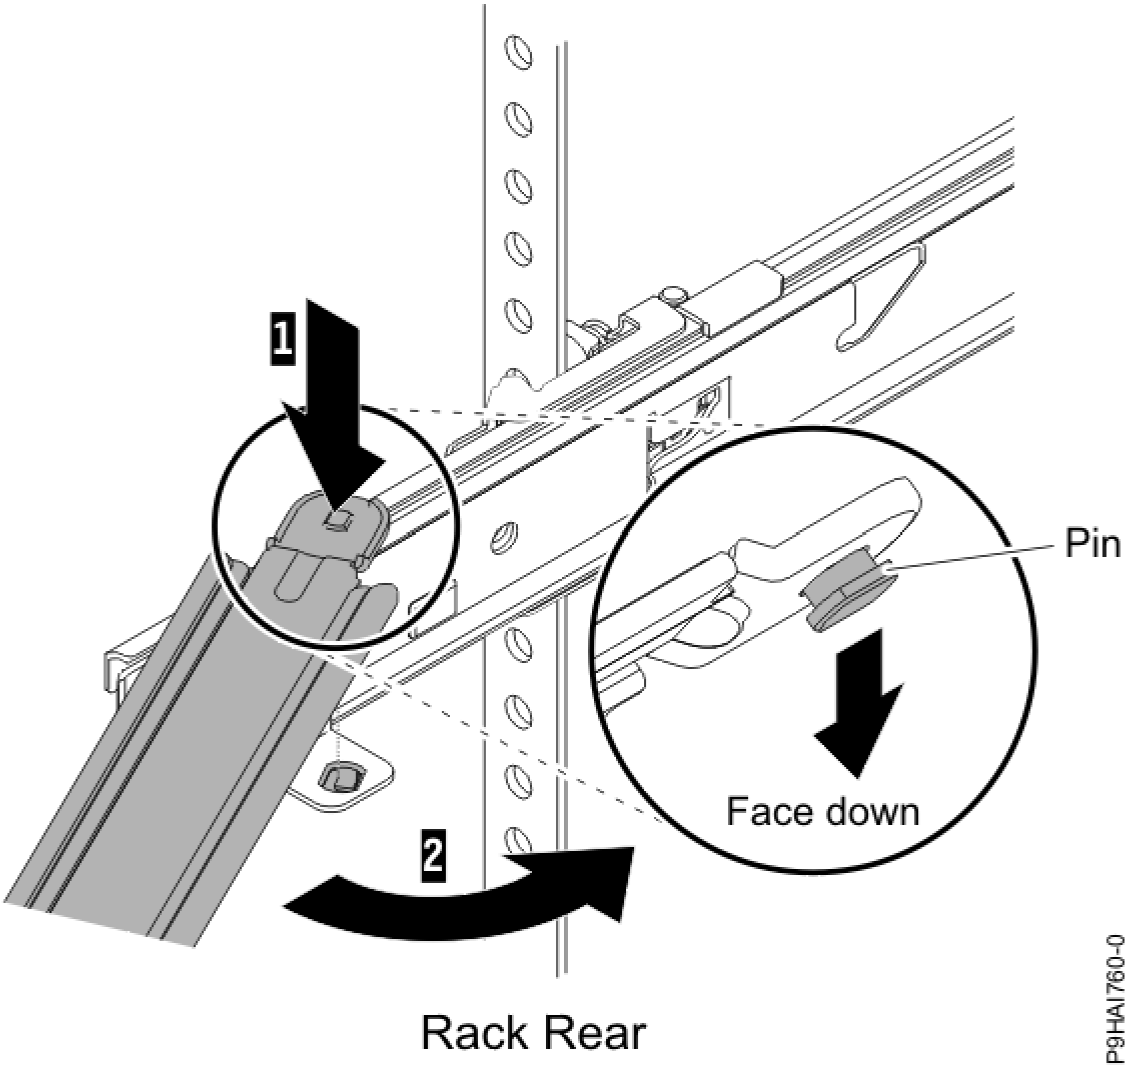

Figure 18. Support arm connection

Note: The cable management support bar must be on top of the slide tab to work correctly.

9. Install the cable management stop bracket (with capital letter **O**) on the unattached end of the support arm. Ensure that the support arm is securely installed.

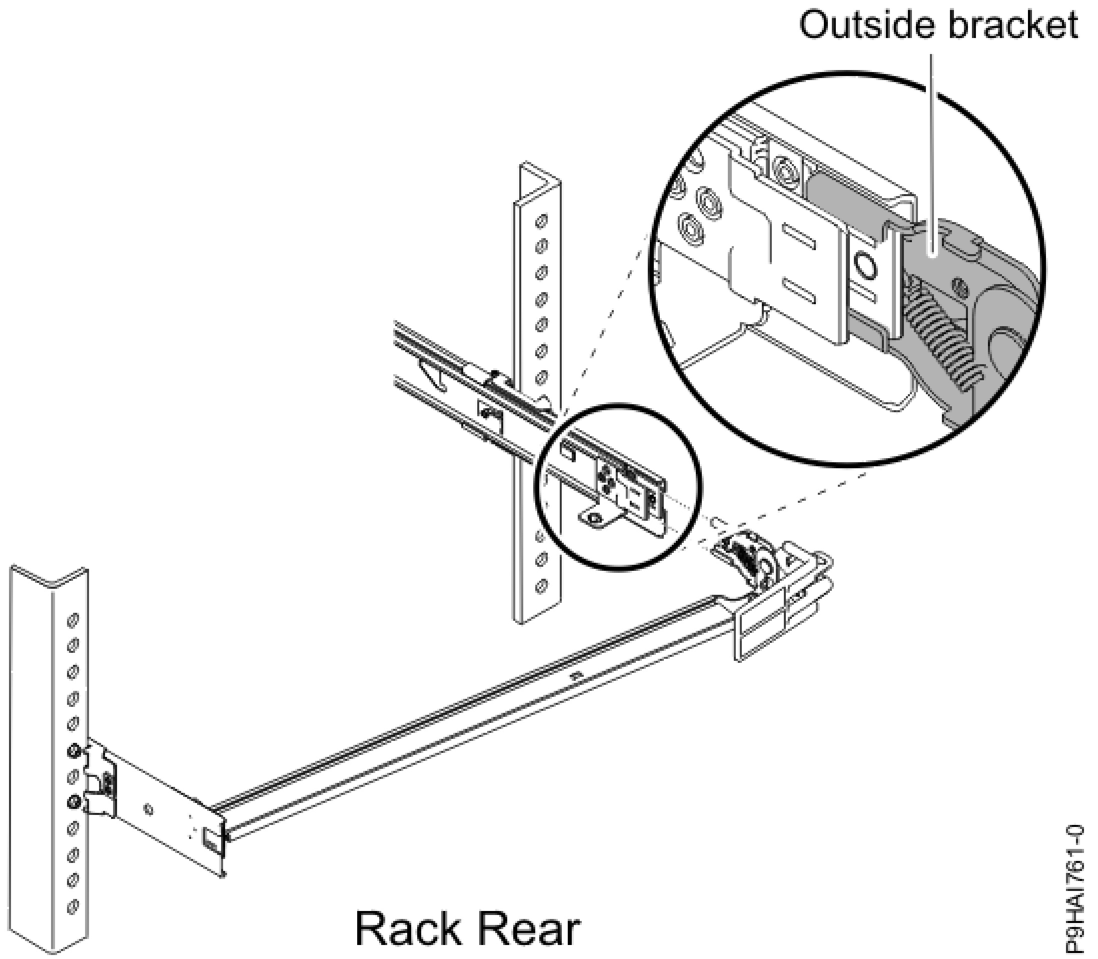

Figure 19. Connecting the stop bracket to the slide rail

Note: The capital letter **O** is marked on cable management arm pins to identify the outside pins.

10. Place the cable-management arm on the support arm. Slide the cable management arm tabs into both the inside and the outside slots of the slide rail. Push the tabs until they snap into place.

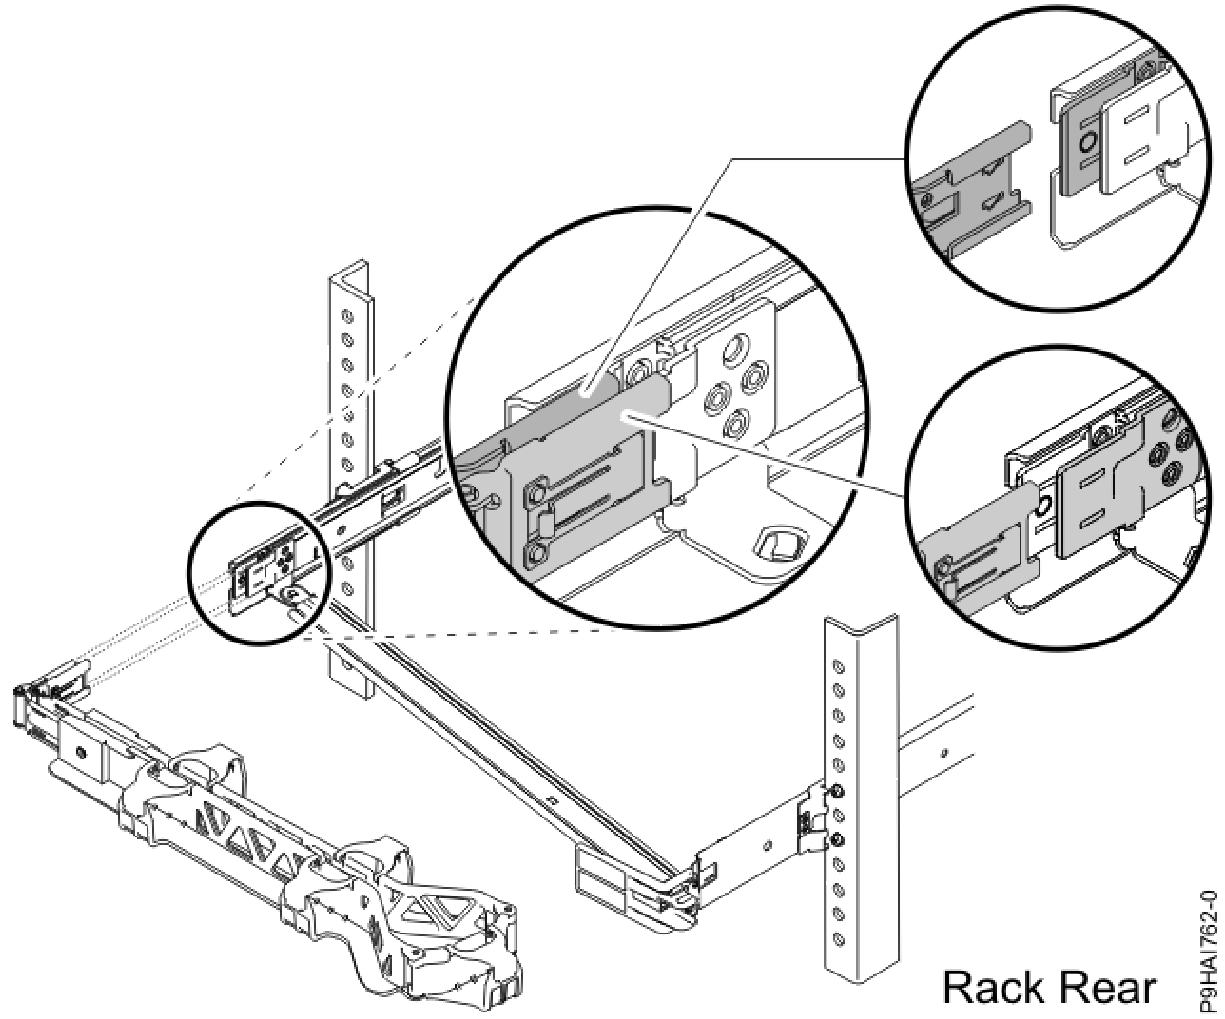

Figure 20. Cable-management arm connection

11. To make it easier to rotate the cable management arm on and off the cable management support arm, you can open the stop bracket by pushing the tabs above and below the cable management support.

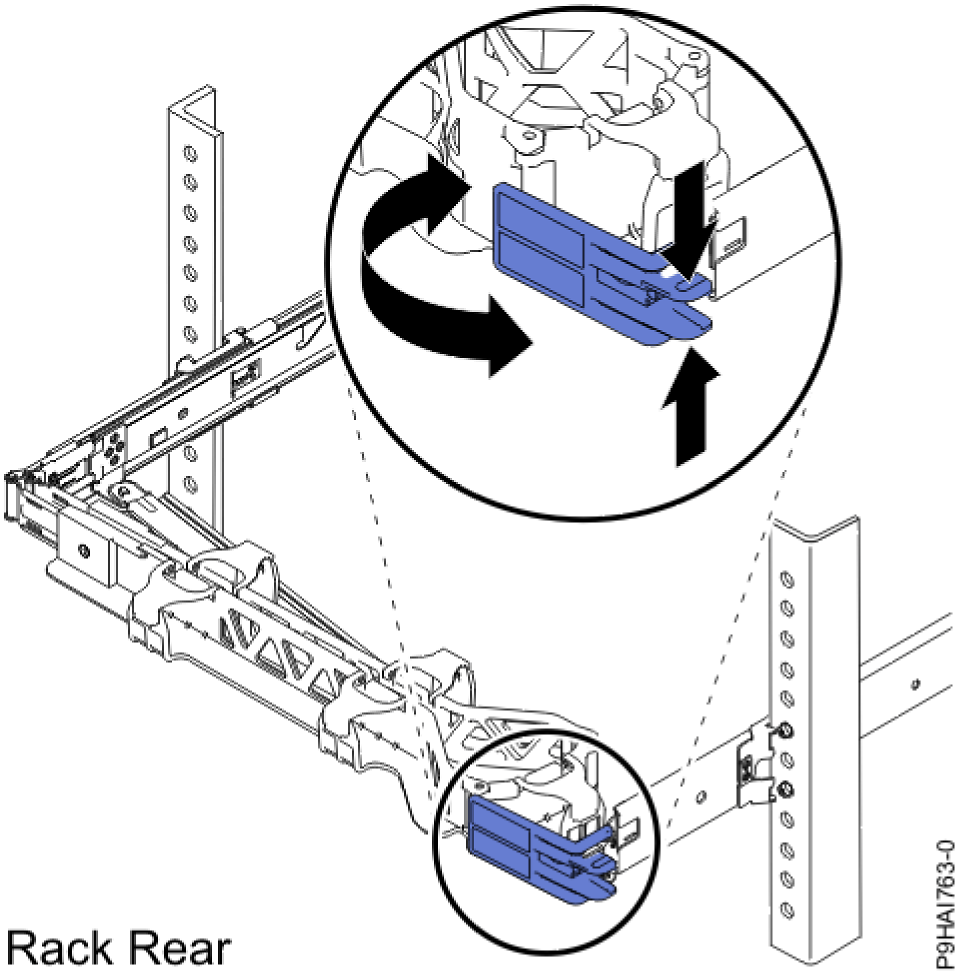

Figure 21. Cable management support stop bracket

12. Attach the power cords and other cables to the rear of the server (including keyboard, monitor, and mouse cables, if required). Route the cables and power cords on the cable-management arm and secure them with cable ties or hook-and-loop fasteners.

    Note: The location of the cable straps might be slightly different in different systems. Use the cable straps that are provided on the rear of the system to retain the cables and prevent them from sagging.

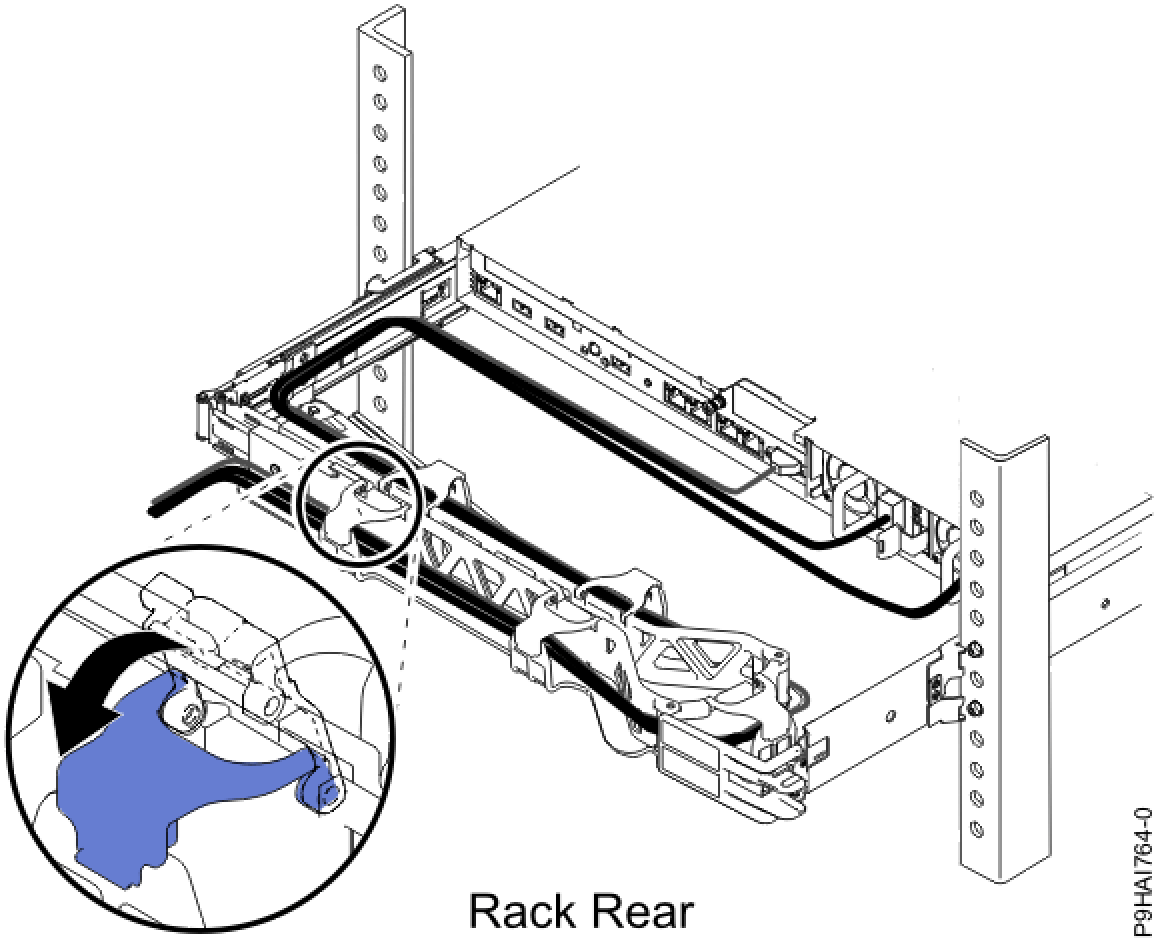

Figure 22. Attaching the Power cord and routing the cable

13. Cables must be bundled with the cable strap for proper movement of the cable management arm.

    Note: Ensure that the cables do not sag below the U space so they do not get caught on the lower systems. Allow slack in all cables to avoid tension in the cables as the cable-management arm moves.

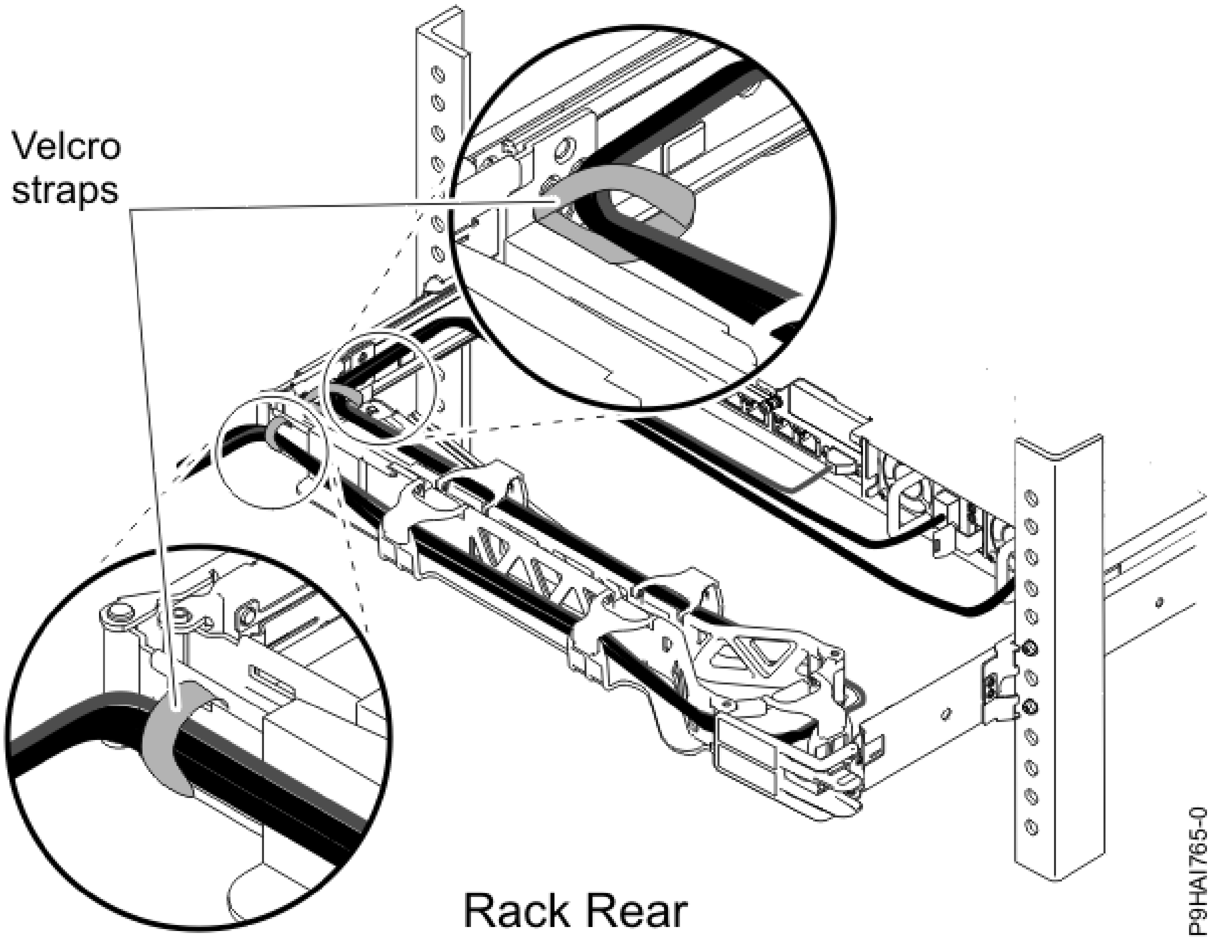

Figure 23. Hook-and-loop fastener

14. If you are shipping the rack with the system installed or if you are in a vibration-prone area, insert the M5 screws into the rear of the slides. Use a cable tie to secure the free end of the cable management arm to the rack if needed.

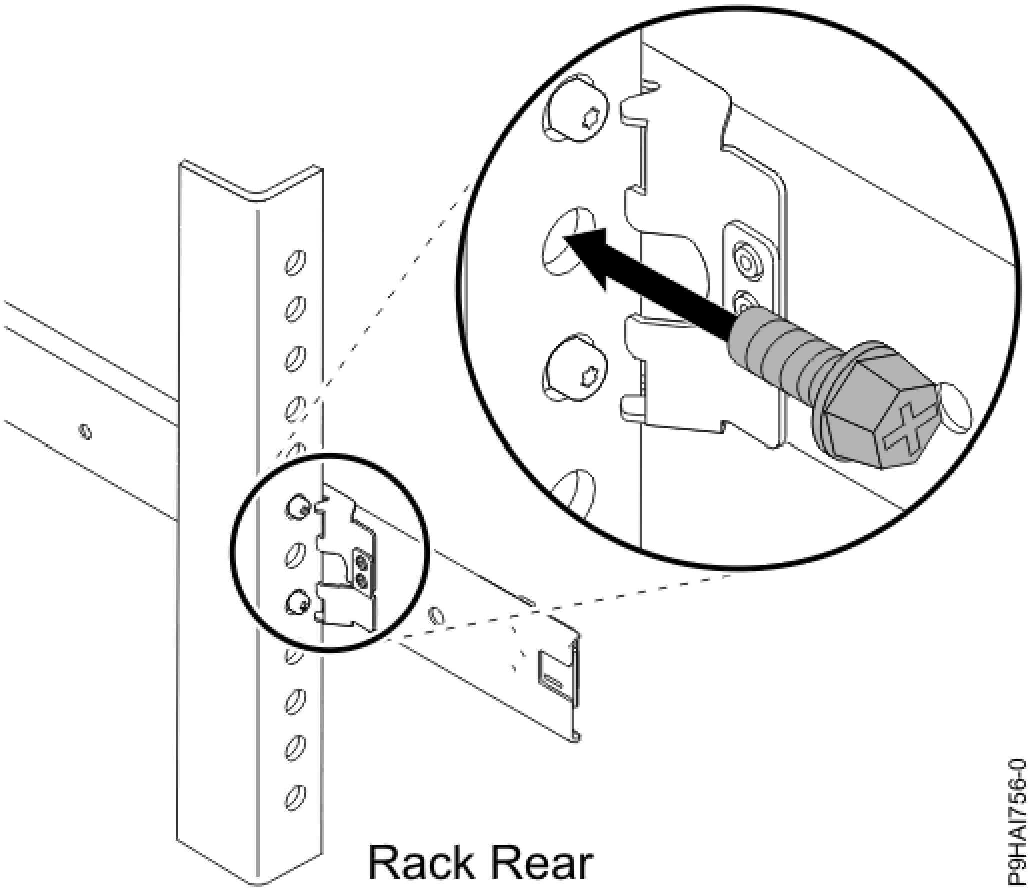

Figure 24. Securing the server for shipping

Installing the HMC virtual appliance

Learn how to install the Hardware Management Console (HMC) virtual appliance.

The HMC virtual appliance can be installed in your existing x86 or POWER® virtualized infrastructure. The HMC virtual appliance supports the following x86 virtualization hypervisors:

- Kernel-based virtual machine (KVM)
- Xen
- VMware

The HMC virtual appliance supports the following POWER virtualization hypervisors:

- PowerVM®

Minimum requirements for running the HMC virtual appliance:

- 16 GB of memory
- 4 virtual processors
- 2 network interfaces (maximum of 4 allowed)
- 1 disk drive that contains 500 GB of available disk space

  Note:

  PowerVM virtualization hypervisor requires 160 GB of disk space. 500 GB of memory is recommended.

  The minimum PowerVM processor requirement is 1.0 processing units and four shared virtual processors in capped sharing mode. Using dedicated processor is not recommended. 16 GB of memory is recommended.

Notes:

1. The processor on the systems that host the HMC virtual appliance must be either an Intel VT-x or an AMD-V hardware virtualization-enabled processor.
2. The HMC virtual appliance DVDs that you receive are not bootable. You must mount the media first and then copy the `.tgz` file from the media. The method to mount the DVD can vary depending on the operating system that you use.
3. The command syntax that are used in the following examples can vary depending on the operating system that you use.

Performance and scale requirements:

- When the HMC is at V9.2.950.0, or later, a single HMC can manage up to 48 systems and 2000 partitions across the systems that are managed by the HMC with the following requirements:
  - 16 GB memory for 1 - 500 partitions
  - 32 GB memory for 500 - 2000 partitions
- You can refer to the corresponding product documentation to find the maximum number of systems and partitions that an HMC can manage, when the HMC is used in combination with Cloud Management Console (CMC), PowerVC, VM High Availability/Disaster Recovery (HA/DR).

Related information

HMC V8 network installation images and installation instructions

Installing the HMC virtual appliance on x86

Learn how to install the Hardware Management Console (HMC) virtual appliance on a x86 environment.

Installing the HMC virtual appliance by using the KVM hypervisor

Learn how to install the Hardware Management Console (HMC) virtual appliance by using the kernel-based virtual machine (KVM) hypervisor.

To install the HMC virtual appliance on KVM, complete the following steps:

Note: The following use the command line interface and require root user authority. The command syntax might vary depending on the operating system.

1. Verify that virtualization packages are installed on systems with Red Hat Enterprise Linux (RHEL) version 7.0 or later.
2. Download the `<KVM vHMC installation filename>.tar.gz` file to the host system.
3. Run the following command: `mkdir -p /var/lib/libvirt/images/vHMC`.
4. Run the following command: `cd /var/lib/libvirt/images/vHMC`.
5. To extract the virtual disk images, run the following command: `tar -zxvf <KVM vHMC installation filename>.tgz`

   Note: In this command, specify the full path of your HMC virtual appliance .tar file.
6. A **domain.xml** file is provided in the `<KVM vHMC installation filename>.tar.gz` file. Complete the following steps:
   a. Edit the **domain.xml** file and verify that the path to your disks is correct. This file contains the string **DISK_PATH**.
   b. Make sure `virtio` is used in the bus value for your disk device.
   c. You can choose to have a different name for your VM. The default name in the **domain.xml** file is **vHMC**.
   d. Verify that the media access control (MAC) address is set in the **domain.xml** file. This file contains the string **MAC_ADDRESS**.

      Note: Remove this line if you want a MAC address to be generated automatically for you.

   e. Verify that your bridges match your Ethernet devices. The default **domain.xml** file specifies one Ethernet.

   f. If you are using the Activation Engine, replace **AEDISK** with the name of Activation Engine virtual disk image. Otherwise, remove the disk element.

7. To define the VM, run the following command: `virsh define <domain>.xml`.
8. To verify that Virtual HMC was added to the list of defined VMs, run the following command: `virsh list --all`.
9. To start the VM, run the following command: `virsh start vHMC`.
10. To determine the Virtual Network Computing (VNC) display number of your console, run the following command: `virsh vncdisplay vHMC`.
11. To connect to your console with a VNC viewer, run the following command: `vncviewer HOSTNAME:ID`(Where ID is the display number, for example `0`).

    Note: If you require remote access, you must drop or configure your firewall to allow access to port 5900.

Installing the HMC virtual appliance by using the Xen hypervisor

Learn how to install the Hardware Management Console (HMC) virtual appliance by using the Xen hypervisor.

The HMC virtual appliance supports Xen version 4.2 or later.

To install the HMC virtual appliance by using the Xen hypervisor, complete the following steps:

Note: The following steps use the command line interface and require root user authority. The command syntax might vary depending on the operating system.

1. Verify that virtualization packages are installed on systems with Red Hat Enterprise Linux (RHEL) version 6.4 or later.
2. Download the `<XEN vHMC installation filename>.tar.gz` file to the host system.
3. Run the following command: `mkdir -p /var/lib/libvirt/images/vHMC`.
4. Run the following command: `cd /var/lib/libvirt/images/vHMC`.
5. To extract the virtual disk images, run the following command: `tar -zxvf <XEN vHMC installation filename>.tgz`

   Note: In this command, specify the full path of your HMC virtual appliance .tar file.
6. A **vhmc.cfg** file is provided in the `<XEN vHMC installation filename>.tar.gz` file. Open the **vhmc.cfg** file in a text editor and edit the following values:

   a. Change the name of the virtual HMC (optional): Edit the **vhmc.cfg** file and verify that the path to your disks is correct. This file contains the string **DISK_PATH**.

   b. Replace **DISK_PATH** with the path for `disk1.img`:

   ```
   disk = [ 'file:DISKPATH,hda,w' ]
   ```

   c. Replace **ethernet adapter** and add MAC address (optional):

   ```
   vif = [ 'type=virtio, model=e1000, bridge=eth0' ]
   ```

   Optional MAC Address:

   ```
   vif = [ 'type=virtio, mac=MACADDRESS, model=e1000, bridge=eth0' ]
   ```

   Note: When the Virtual HMC is rebooted, the Xen hypervisor automatically regenerates a MAC address. Adding the optional MAC Address solves this issue.

   d. Replace **FLOPPYPATH** (if you are using the Activation Engine):

   ```
   device_model_args = [ "-fda", "FLOPPYPATH" ]
   ```

7. To create and start the VM, run the following command: `xl create vHMC.cfg`.
8. To check that the VM was added to the list of defined virtual machines, run the following command: `xl list`.
9. To access the VM local console, run the following command: `vncviewer localhost 0`.

Installing the HMC virtual appliance by using VMware ESXi

Learn how to install the Hardware Management Console (HMC) virtual appliance by using VMware ESXi.

You can install the HMC virtual appliance on VMware ESXi by using the graphical user interface on the vSphere client to deploy the Open Virtualization Format (OVF) template.

Note: You can install the HMC virtual appliance on VMware ESXi version 6.0 or later.

To install the HMC virtual appliance on VMware ESXi by using the vSphere client, complete the following steps:

Note: The command syntax might vary depending on the operating system.

1. Obtain the Tar archive file: `<VMware vHMC installation file name>.tgz`.
2. Use the `tar` command to extract the OVA file from the Tar archive file.
3. Start the vSphere client and log in to the ESXi host.
4. From the **File** menu, select **Deploy OVF template**.
5. Click **Browse** and select the OVA file.
6. Click **Next**.
7. After the deployment is completed, click **Close** and select the HMC virtual appliance icon to power the HMC virtual appliance on.

Installing the HMC virtual appliance on POWER

Learn how to install the Hardware Management Console (HMC) virtual appliance on a virtualized POWER environment.

Installing the HMC virtual appliance on PowerVM (logical partition)

Learn how to install the Hardware Management Console (HMC) virtual appliance on a PowerVM environment.

The HMC virtual appliance supports POWER9 servers on firmware level FW910 or later. For more information, see Supported Linux distributions for POWER8® and POWER9 Linux on Power systems (https://www.ibm.com/support/knowledgecenter/en/linuxonibm/liaam/liaamdistros.htm).

Notes:

1. You cannot manage the server that hosts the HMC virtual appliance.
2. You cannot manage the server that hosts another HMC virtual appliance which is managing the server that hosts this HMC virtual appliance.

   For example, HMC virtual appliance A is running on server A and HMC virtual appliance B is running on server B. HMC virtual appliance A cannot manage server B and HMC virtual appliance B cannot manage server A at the same time. One of the HMC virtual appliance can manage the other server, but both HMC virtual appliance cannot manage each other at the same time.

.

Create automated HMC installation image (optional)

You can create an automated HMC installation image that automatically installs the HMC virtual appliance without prompting for the **HMC Installallation** wizard.

Note: The HMC virtual appliance on PowerVM does not provide graphics adapter support for adapters that are assigned to the partition. You can use a supported web browser to connect to the HMC for user interface support.

To create an automated HMC installation image, complete the following steps:

1. Create two directories by running the following commands: `mkdir -p oldiso` and `mkdir -p newiso`.
2. Mount the HMC installation image to the **oldiso** directory by running the following command: `sudo mount -o loop <image_path> oldiso`.
3. Copy the contents of the **oldiso** directory to the **newiso** directory by running the following command: `cp -r oldiso/* newiso`.
4. If the Grub configuration file (**newiso/boot/grub/grub.cfg**) has 2 lines that start with **menuentry**, then remove the first **menuentry** section by running the following command: `sed -i '/\"Install Hardware Management Console\"/,+4d' newiso/boot/grub/grub.cfg`.
5. Edit the Grub file for the automated install by running the following command: `sed -i 's/biosdevname=0/biosdevname=0 mode=auto optype=Install/' newiso/boot/grub/grub.cfg`.
6. Make the Grub file read-only by running the following command: `sudo chown 0444 newiso/boot/grub/grub.cfg`.
7. Create a new HMC installation ISO by running the following command: `mkisofs -o <new_iso_name> -V <ISO label> -f -r -T -udf --allow-limited-size --netatalk -chrp-boot -iso-level 4 -part -no-desktop -quiet newiso` (where **ISO label** must be `HMC-<hmc version release number>`, for example `HMC-8.0.870.0`).

Note: For more information about setting up the Activation Engine and the configuration file, see "Using the Activation Engine for the HMC virtual appliance" on page 45.

Logical volume setup

To set up the logical volume, complete the following steps:

1. Select a managed system.
2. From the menu pod, select **System Actions** > **Power VM** > **Virtual Storage**.
3. Select **Manage System VIOS** > **Action** > **Manage Virtual Storage**.
4. Select the **Virtual Disks** tab.
5. Click **Create virtual disk** and enter the following information:
   - **Virtual disk name**: The name of the virtual disk.
   - **Storage pool name**: The name of the storage pool.
   - **Virtual disk size**: The size of the virtual disk.
   - **Assigned partition**: The name of the logical partition.

   Note: A minimum of 160 GB disk space is required (500 GB disk space is recommended).

Installation media setup - create media library

To create a media library, complete the following steps:

1. Select a managed system.
2. From the menu pod, select **System Actions** > **Power VM** > **Virtual Storage**.
3. Select **Manage System VIOS** > **Action** > **Manage Virtual Storage**.
4. Select the **Optical Devices** tab.
5. Click **Create Library** and enter the following information:
   - **Storage pool**: The name of the storage pool.

- **Media library size**: The size of the media library.

6. Click **OK**.

Installation media setup - upload media to VIOS

To upload media to Virtual I/O Server (VIOS), complete the following steps:

1. Log in to VIOS.
2. In VIOS root mode, run the following command: `oem_setup_env`.
3. To allow NFS connection, run the following command: `nfso -o nfs_use_reserved_ports=1`.
4. To mount the NFS into the local VIOS folder, run the following command: `mount <server_ip>:/Mountpoint <local_folder>`.
5. To verify that the NFS mount includes your HMC installation ISO and Activation Engine configuration image (optional), run the following command: `ls`.

Installation media setup - link media to media library

To link media to the media library, complete the following steps:

1. Navigate back to **Manage System VIOS** > **Action** > **Manage Virtual Storage** and select the **Optical Devices** tab.
2. From the **Virtual Optical Media** section, select **Add Media** from the **Actions** menu.
3. From the **Add Virtual Media** window, select **Add existing file from VIOS filesystem** and enter the following information:
   - **Media name**: The name of the media (for example, `HMCInstall` or `AEDrive`).
   - **Optical media file name**: The file name of the installation ISO file (for example, `01234567-ppc64ie.iso`).
4. Click **OK**.
5. If you created an Activation Engine configuration image, repeat steps 3 - 4 to add the Activation Engine configuration image. Otherwise, continue to step 6.
6. Verify that the optical media is uploaded to the media library by verifying that the media name is shown in available **Virtual Optical Media** list.

Logical partition setup

To set up the logical partition, complete the following steps:

1. Select a managed system.
2. From the menu pod, select **System Actions** > **Partitions** > **Partitions**.
3. Click **Create Partition** and enter the following information:
   - **Parititon Name**: The name of the partition.
   - **Partition ID**: The ID of the partition.
   - **Partition Type**: Select the operating system (**AIX/Linux** or **IBM i**).
4. Click **OK**.
5. Allocate the number of processors and the amount of memory for the partition.

   Note: A minimum of four virtual processors and 8 GB of memory is required.
6. From the menu pod, select **Partition Actions** > **Virtual I/O** > **Virtual Networks**.
7. Click **Attach Virtual Network** and select the **Show and attach new virtual ethernet adapters** check box. From the table, select the virtual network adapters that you want to attach to the logical partition.

   Note: A maximum of four virtual network adapters is allowed.

8. From the menu pod, select **Partition Actions** > **Virtual I/O** > **Virtual Storage**.
9. From the **Virtual Optical Device** tab, click **Add Virtual Optical Device**.
10. Enter the **Device Name** (for example, `HMCInstall` or `AEDrive`) and select the wanted Virtual I/O Server from the table.

    Note: Installing the `AEDrive` is optional.
11. Click **OK**.
12. Verify that the virtual optical devices that you added from step 10 is now listed in the table.
13. From the **Action** menu, click **Load**.
14. Select the media file to assign to the logical partition and click **OK**.
15. Verify that the virtual optical devices that you loaded from step 13 is now listed in the table.

Starting the HMC virtual appliance

Note: When you install the HMC virtual appliance on a partition by using the HMC ISO image file, you will not have local graphical console access to the web user interface.

To start the HMC virtual appliance on PowerVM, complete the following steps:

1. Select the managed partition.
2. Open an active connection to the logical partition by selecting **Actions** > **Console** > **Open Terminal Window**.
3. Activate the logical partition by selecting **Actions** > **Activate**.
4. Select **Activate (Normal)** and **Current Configuration**.
5. Click **Finish**.
6. Switch to the terminal window.
7. From the **Boot** menu, select **1 = SMS Menu**.
8. From the **Main** menu, select **5 = Select Boot Options**.
9. From the **Multiboot** menu, select **1 = Select Install/Boot Device**.
10. From the **Select Device Type** menu, select **5 = List all devices**.
11. Select the `HMCInstall` device based on the device location.
12. Select **2. Normal Mode Boot**.
13. Select **1. Yes** to confirm.
14. Follow the onscreen instructions from the **HMC Install** wizard.

    Note: Skip this step if you used an automated HMC installation image.
15. After the installation completes and the system starts, you must select a language from the **language selection** dialog box.
16. Accept the license agreement.

    Note: Ensure that the command controller is ready to accept commands before you run any commands. For example, running the **`lshmc -V`** command until it succeeds.
17. Log in as `hscroot` and use the **chhmc** command to configure the network.

    The following example shows the sequence of **chhmc** commands that can be used to configure the network and enable Secure Shell (SSH) and remote web access on the HMC.

    ```
    chhmc -c network -s modify -i ethX -a <hmc ip address> -nm <hmc network mask> --lparcomm on
    chhmc -c network -s modify -h <hmc hostname> -d <hmc domain name> -g <gateway ip>
    chhmc -c network -s add -ns <name server> -ds <domain search>
    chhmc -c ssh -s enable
    chhmc -c ssh.name -s add -a <ip address>
    chhmc -c SecureRemoteAccess.name -s add -a <ip address>
    chhmc -c remotewebui -s enable -i ethX
    hmcshutdown -r -t now
    ```

- **ethX** is the network interface name to configure.
- **hmc ip address** is the IP address of your HMC.
- **hmc network mask** is the network mask of your HMC.
- **hmc hostname** is the host name of your HMC.
- **hmc domain name** is the domain name of your HMC.
- **gateway ip** is the IP address of the gateway on your network.
- **name server** is the name server address of your network.
- **domain search** is the names of the domains that you want the HMC to search on.
- To allow access on all IP addresses, use **-a 0.0.0.0 -nm 0** in place of **ip address**.

Note: When you use multiple virtual Ethernet adapters, run the command **`cat /etc/sysconfig/network-scripts/ifcfg-ethX`** on the HMC virtual appliance on each interface. Compare the media access control (MAC) address against what the HMC shows in the adapter view of the virtual network of the partition. You can click **View Virtual Ethernet Adapter Settings** for more information on the virtual Ethernet adapters. This step helps you determine the correct interface to use.

18. Restart the system.

Using the Activation Engine for the HMC virtual appliance

Learn how to use the Activation Engine for the Hardware Management Console (HMC) virtual appliance.

Activation Engine is a framework that allows various components within a virtual machine to be configured during system startup. To use the Activation Engine, you need to set up an XML configuration profile to allow the HMC virtual appliance to be in a ready-to-manage state on first start. For more information about configuring the XML configuration profile, see "Setting up the configuration profile for the Activation Engine" on page 46. The configuration file can be used to configure the following options:

- Set Default Keyboard (US)
- Default Locale (US)
- Disable Keyboard Setup
- Disable Display Setup
- License Agreement and Machine Code Agreement
- Disable Setup Wizard
- Disable Call Home Wizard
- Configure up to four Network Interface Cards
- Configure Firewall Settings for each Interface
- Configure Network interface as IPv4 DHCP Server
- Configure Private and Open Interface
- Configure Default Gateway Interface Device

Note: The number of Ethernet adapters that is defined in the **vHMC-Conf.xml** configuration file must correlate with the defined Network adapters in the **domain.xml**, **vHMC.cfg**, or **VMWare** configuration file.

The Activation Engine requires a virtual disk that holds an XML configuration. You can edit the **user_data** file with a text editor and use the XML configuration guide that is shown in the following example.

To create a virtual ISO disk image with Activation Engine configuration in a Linux environment, complete the following steps:

1. Create a directory:

```
mkdir -p config-drive/openstack/latest
```

2. Copy the edited **user_data** file into the directory:

   ```
   cp user_data config-drive/openstack/latest
   ```

3. Create a virtual disk image with the Activation Engine configuration:

   ```
   mkisofs -R -V config-2 -o AEdrive.iso config-drive
   ```

Setting up the configuration profile for the Activation Engine

Learn how to set up the Activation Engine configuration file by using XML tags.

Configuration file

Use the following example of the configuration file to learn about the XML tags.

```
<vHMC-Configuration>
   <ConfigurationVersion>2.0</ConfigurationVersion>
   <LicenseAgreement></LicenseAgreement>
   <AcceptLicense>Yes</AcceptLicense>
   <Locale>en_US.UTF-8</Locale>
   <SetupWizard>No</SetupWizard>
   <SetupCallHomeWizard>No</SetupCallHomeWizard>
   <SetupKeyboard>No</SetupKeyboard>
   <SetupDisplay>No</SetupDisplay>
   <Ethernet Enable='Yes' DefaultGatewayDevice='Yes' PrivateInterface='No'>
      <Hostname></Hostname>
      <Domain></Domain>
      <DNSServers></DNSServers>
      <IPV4Config>
         <NetworkType></NetworkType>
         <IPAddress></IPAddress>
         <Netmask></Netmask>
         <Gateway></Gateway>
      </IPV4Config>
      <IPV6Config>
         <NetworkType></NetworkType>
         <IPAddress></IPAddress>
         <Gateway></Gateway>
      <IPV6Config>
      <Firewall>
         <PEGASUS>Enabled</PEGASUS>
         <RPD>Enabled</RPD>
         <FCS>Enabled</FCS>
         <I5250>Enabled</I5250>
         <PING>Enabled</PING>
         <L2TP>Disabled</L2TP>
         <SLP>Enabled</SLP>
         <RSCT>Enabled</RSCT>
         <SECUREREMOTEACCESS>Enabled</SECUREREMOTEACCESS>
         <SSH>Enabled</SSH>
         <NTP>Disabled</NTP>
         <SNMPTraps>Disabled</SNMPTraps>
         <SNMPAgents>Disabled</SNMPAgents>
      </Firewall>
   </Ethernet>
   <NTPServers>
      <ntpparam ntpserver="" ntpversion=""/>
   </NTPServers>
</vHMC-Configuration>
```

XML tags for the configuration file

XML tags are used in the Activation Engine configuration file to set specific values for various attributes. You can manually set these values in the Activation Engine configuration file. Use the following table to see a description of each tag and its allowed values:

Table 9. XML tags

Tags	Description	Acceptable values	Notes
ConfigurationVersion	Required element that defines the configuration version to use.	**2.0**	
LicenseAgreement	Required element that displays the HMC virtual appliance license agreement.		
AcceptLicense	Required element to accept the HMC virtual appliance license agreement.	• **Yes**: Accepts the HMC license agreement. • **No**: Prompts User to Accept HMC License Agreement	If an invalid value is entered, the Activation Engine uses the default setting of **No**.
Locale	Required element to define locale settings.	**en_US.UTF-8**	If an invalid value is entered, the Activation Engine uses the default setting of **US**.
SetupWizard	Required element to enable or disable the **HMC Setup** wizard.	• **Yes**: Displays the **HMC Setup** wizard. • **No**: Disables the **HMC Setup** wizard display.	If an invalid value is entered, the Activation Engine uses the default setting of **Yes**.
SetupCallHomeWizard	Required element to enable or disable the **HMC Call Home** wizard.	• **Yes**: Displays the **HMC Call Home** wizard. • **No**: Disables the **HMC Call Home** wizard display.	If an invalid value is entered, the Activation Engine uses the default setting of **Yes**.
SetupKeyboard	Required element to define the keyboard configuration.	• **Yes**: Prompts the user for keyboard configuration. • **No**: Accepts default keyboard configuration (US).	If an invalid value is entered, the Activation Engine uses the default setting of **Yes**.
SetupDisplay	Required element to enable or disable the display configuration.	• **Yes**: Prompts the user for display configuration. • **No**: Accepts default display configuration.	If an invalid value is entered, the Activation Engine uses the default setting of **Yes**.

Table 9. XML tags (continued)

Tags	Description	Acceptable values	Notes
Ethernet	Required element that holds values for Ethernet adapter configurations. A maximum of four Ethernet adapters can be configured.	**Enable**: • **Yes**: Configure this adapter. • **No**: Do not configure this adapter. **DefaultGatewayDevice**: • **Yes**: Configure this adapter as the main network adapter. • **No**: Do not configure this adapter as the main network adapter. **PrivateInterface**: • **Yes**: Configure this adapter as a private interface. **Yes** is required to configure interface as an IPv4 DHCP Server. • **No**: Do not configure this adapter as a private interface. **No** is required to configure interface as IPv4 static type.	The Activation Engine runs the default configuration if any invalid values are entered within the Ethernet adapter section or if multiple **Default Gateway Devices** are defined. Optional elements can be omitted from the configuration. At least one IPV4 or IPV6 configuration is required. If you do not specify an IP configuration, the Activation Engine uses the default configuration.
HostName	Optional element to define the network host name.	Any valid host name string.	If the element is not defined, the Activation Engine uses the default local host **HostName** value.
Domain	Optional element to define the network domain.	Any valid domain value (for example, **example.us.com**).	If the element is not defined, the Activation Engine uses the default empty **Domain** value.
DNSServers	Optional element to define the network DNS servers.	It is acceptable to have one DNS Server value or up to three valid IPv4 or IPv6 addresses that are separated by a comma. • Example 1: IPv4: 8.3.2.1 IPv6: 2001:4860:4860::8888 • Example 2: IPv4: 8.3.2.1,8.5.4.1 IPv6: 2001:4860:4860::8888,2001:4860:4860::8844 • Example 3: IPv4: 8.3.2.1,8.5.4.1,8.4.3.2 IPv6: 2001:4860:4860::8888,2001:4860:4860::8844, ::ffff:903:201	If the element is not defined, the Activation Engine uses the default empty **DNSServers** value.

Table 9. XML tags (continued)

Tags	Description	Acceptable values	Notes
IP4Config	Optional element to define IPv4 configuration settings.	**IPType**: Required element to define IPv4 configuration type. • **Static**: Configure this adapter by using static configuration. • **DHCP**: Configure this adapter by using DHCP configuration. • **DHCPServer**: Configure this adapter to be IPv4 DHCP server (requires **PrivateInterface** to be set to **Yes**). **IPAddress**: Optional element that is required only if **Static** or **DHCPServer** configuration is selected. • **Static Configuration**: Any valid IPv4 address value. • **DHCPServer Configuration**: Any DHCP server IP within the IP range. **Netmask**: Optional element that is required only if **Static** configuration is selected. • Any valid IPv4 netmask value. **Gateway**: Optional element that is required only if **Static** configuration is selected. • Any valid IPv4 netmask value.	
IP6Config	Optional element to define IPv6 configuration settings.	**IPType**: Required element to define IPv6 configuration type. • **Static**: Configure this adapter by using static configuration. • **DHCP**: Configure this adapter by using DHCP configuration. **IPAddress**: It is acceptable to have long or short form IPv6 format and long or short form IPv6 prefix. • Example 1: IPv6: 2001:4860:4860:0000:0000:0000:0000:8888 • Example 2: IPv6: 2001:4860:4860::8888 • Example 3: IPv6: 2001:4860:4860::8888/128 If no prefix is specified, the Activation Engine uses the default setting of /64 prefix. **Gateway**: • Any valid IPv6 address value.	

Table 9. XML tags (continued)

Tags	Description	Acceptable values	Notes
Firewall	Optional element to define firewall settings.	**PEGASUS**: • **Enabled**: Allows the PEGASUS ports to be open. • **Disabled**: Disables PEGASUS ports. **RPD**: • **Enabled**: Allows the RMC ports to be open. • **Disabled**: Disables RMC ports. **FCS**: • **Enabled**: Allows the FCS ports to be open. • **Disabled**: Disables FCS ports. **I5250**: • **Enabled**: Allows the 5250 ports to be open. • **Disabled**: Disables 5250 ports. **PING**: • **Enabled**: Allows the Ping port to be open. • **Disabled**: Disables Ping port. **L2TP**: • **Enabled**: Allows the L2TP ports to be open. • **Disabled**: Disables L2TP ports. **SLP**: • **Enabled**: Allows the SLP ports to be open. • **Disabled**: Disables SLP ports. **RSCT**: • **Enabled**: Allows the RSCT ports to be open. • **Disabled**: Disables RSCT ports. **SECUREREMOTEACCESS**: • **Enabled**: Allows the secure remote access ports to be open. • **Disabled**: Disables secure remote access ports. **SSH**: • **Enabled**: Allows the SSH port to be open. • **Disabled**: Disables SSH port.	

Table 9. XML tags (continued)

Tags	Description	Acceptable values	Notes
Firewall	Optional element to define firewall settings.	**NTP**: • **Enabled**: Allows the NTP ports to be open. • **Disabled**: Disables NTP ports. **SMNPTraps**: • **Enabled**: Allows the SMNP traps ports to be open. • **Disabled**: Disables SMNP traps ports. **SMNPAgents**: • **Enabled**: Allows the SMNP agents ports to be open. • **Disabled**: Disables SMNP agents ports.	

NTPServers	The **NTPServers** tag is needed if you want to configure up to five NTP servers within a HMC virtual appliance.	**NTPServers**: Accepts `<ntpparam ntpserver="server" ntpversion="version"/>`
		ntpparam:
		• **ntpserver**: Accepts any valid IPv4 or IPv6 values and valid host names.
		• **ntpversion**: Accepts 1-4 numeric value
		Example:
		`<NTPServers>`
		`  <ntpparam ntpserver=`
		`   "test.austin.ibm.com"`
		`   ntpversion="2"/>`
		`  <ntpparam`
		`ntpserver="192.168.34.1"`
		`   ntpversion="4"/>`
		`  <ntpparam`
		`ntpserver="::ffff:903:201"`
		`` ntpversion="3"/>` ``
		`</NTPServers>`	

Configuring the HMC

Learn how to set up your network connections, configure your HMC, complete postconfiguration steps, and upgrade and update your HMC.

Choosing network settings on the HMC

Learn about the network settings that you can use on the Hardware Management Console (HMC).

HMC network connections

Learn how the Hardware Management Console HMC can be used in a network.

You can use different types of network connections to connect your HMC to managed systems. For more information about how to configure the HMC to connect to a network, see "Configuring the HMC" on page 67. For more information about using the HMC on a network, see the following information:

Types of HMC network connections

Learn how to use the HMC remote management and service functions by using your network.

The HMC supports the following types of logical communications:

HMC to managed system

Used to perform most of the hardware management functions, in which HMC issues control function requests through the service processor of the managed system. The connection between the HMC and the service processor is sometimes referred to as the *service network*. This connection is required for managed system management.

HMC to logical partition

Used to collect platform-related information (hardware error events, hardware inventory) from the operating systems that are running on logical partitions, and to coordinate certain platform activities (dynamic LPAR, concurrent repair) with those operating systems. If you want to use service and error notification features, you must create this connection.

HMC to BMC

Note: The baseboard management controller (BMC) connection is applicable only to HMC model 7063-CR1.

Used to perform service and maintenance tasks. The BMC connection is used to load and maintain the HMC firmware on the system. This connection is required for access to the BMC on the HMC.

HMC to remote users

Provides remote users with access to HMC functions. Remote users can access the HMC in the following ways:

- By using the web browser to access all the HMC GUI functions remotely.
- By using Secure Socket Shell (SSH) to access the HMC command line functions remotely.
- By using a virtual terminal server for remote access to virtual logical partition consoles.

HMC to service and support

Used to transmit data, such as hardware error reports, inventory data, and microcode updates, to and from your service provider. You can use this communications path to make automatic service calls.

Your HMC can support up to four separate physical Ethernet interfaces, depending on the model. The stand-alone version of the HMC supports only three HMC interfaces, by using one integrated Ethernet adapter and up to two plug-in adapters. Use each of these interfaces in the following ways:

- One network interface can be used exclusively for HMC-to-managed system communications, which means that only the HMC and service processors of the managed systems would be on that network. One or more network interfaces can be used exclusively for HMC-to-managed system communications, which means that only the HMC and service processors of the managed systems are on that network. Even though the network interfaces into the service processors are encrypted for the Secure Sockets Layer (SSL) Protocol and password-protected, having a separate dedicated network can provide a higher level of security for these interfaces.
- An open network interface would typically be used for the network connection between the HMC and the logical partitions on the managed systems, for the HMC-to-logical partition communications. You can also use this open network interface to manage the HMC remotely.
- Optionally, you can use a third interface to connect to logical partitions and manage the HMC remotely. This interface can also be used as a separate HMC connection to different groups of logical partitions. For example, you might want to have an administrative LAN that is separate from the LAN on which all the usual business transactions are running. Remote administrators can access the HMC and other managed units by using this method. Sometimes the logical partitions are in different Network security domains, perhaps behind a firewall, and you might want to have different HMC network connections into each of those two domains.

Web browser requirements for HMC

The Hardware Management Console (HMC) version 9.1.0 is supported by Google Chrome version 57, Microsoft Internet Explorer (IE) version 11.0, Mozilla Firefox versions 45 and 52 Extended Support Release (ESR), and Safari version 10.1.

If your browser is configured to use an Internet proxy, a local IP addresses should be included in the exception list. Consult your network administrator for more information on the exception list. If you still need to use the proxy to get to the HMC, enable Use HTTP 1.1 through proxy connections under the Advanced tab in your Internet Options window.

Session cookies need to be enabled in order for ASMI to work when connected to HMC remotely. The asm proxy code saves session information and uses it. Follow the steps to enable the session cookies.

Enabling session cookies in Internet Explorer.

1. Select Tools and Click Internet Options
2. Select Privacy and Click Advanced
3. Ensure that the Always allow session cookies is checked. If not, select the Override automatic cookie handling and select Always allow session cookies.
4. Select Prompt under First-party Cookies and Third-party Cookies
5. Click OK.

Enabling session cookies in Firefox.

1. Select Tools and click Options
2. Click Cookies
3. Select Allow sites to set cookies.
4. Select Exceptions and add HMC.
5. Click OK.

Private and open networks in the HMC environment
The Hardware Management Console (HMC) can be configured to use open and private networks. Private networks allow the use of a selected range of nonroutable IP addresses. A *public*, or "open" network describes a network connection between the HMC to any logical partitions and to other systems on your regular network.

Private networks

The only devices on the HMC private network are the HMC itself and each of the managed systems to which that HMC is connected. The HMC is connected to each managed system's Flexible Service Processor (FSP).

On most systems, the FSP provides two Ethernet ports that are labeled **HMC1** and **HMC2**. You to connect up to two HMCs.

Some systems have a dual-FSP option. In this situation, the second FSP acts as a redundant backup. The basic setup requirements for a system with two FSPs are essentially the same as a system without a second FSP. The HMC must be connected to each FSP, so more network hardware is required (for example, a LAN switch or hub) when there is more than one FSP or multiple managed systems.

Note: Each FSP port on the managed system must be connected to only one HMC.

Public networks

The open network can be connected to a firewall or router for connecting to the internet. Connecting to the internet allows the HMC to call home when any hardware errors need to be reported.

The HMC itself provides its own firewall on each of its network interfaces. A basic firewall is automatically configured when you run the HMC Guided Setup wizard, but you customize your firewall settings after the initial HMC installation and configuration.

HMC as a DHCP server
You can use the Hardware Management Console (HMC) as a Dynamic Host Configuration Protocol (DHCP) server.

If you want to configure the first network interface as a private network, you can select from a range of IP addresses for the DHCP server to assign to its clients. The selectable address ranges include segments from the standard nonroutable IP address ranges.

In addition to these standard ranges, a special range of IP addresses is reserved for IP addresses. This special range can be used to avoid conflicts in cases where the HMC attached open networks are using one of the nonroutable address ranges. Based on the range that is selected, the HMC network interface on the private network is automatically assigned the first IP address of that range, and the service processors are then assigned addresses from the rest of the range.

The DHCP server in the HMC uses automatic allocation, which means that each unique service processor Ethernet interface is reassigned the same IP address each time it is started. Each Ethernet interface has a unique identifier that is based on a built-in Media Access Control (MAC) address, which allows the DHCP server to reassign the same IP parameters. You can configure both **eth0** and **eth1** HMC ports to serve DHCP addresses.You can configure both **eth0** and **eth1** HMC ports to serve DHCP addresses.

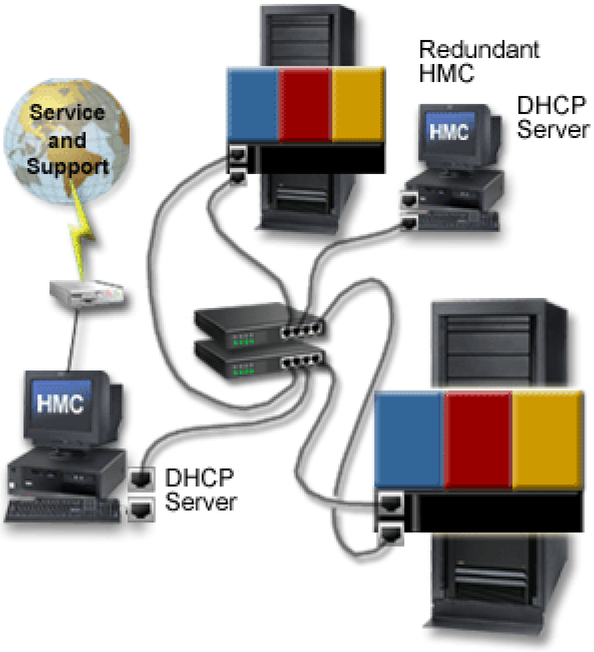

Figure 25. Private network with one HMC as a DHCP server

Note: If you are using IPv6, the discovery process must be done manually. For IPv6, automatic discovery is not available.

For more information about how to configure the HMC as a DHCP server, see "Configuring the HMC as a DHCP server" on page 75.

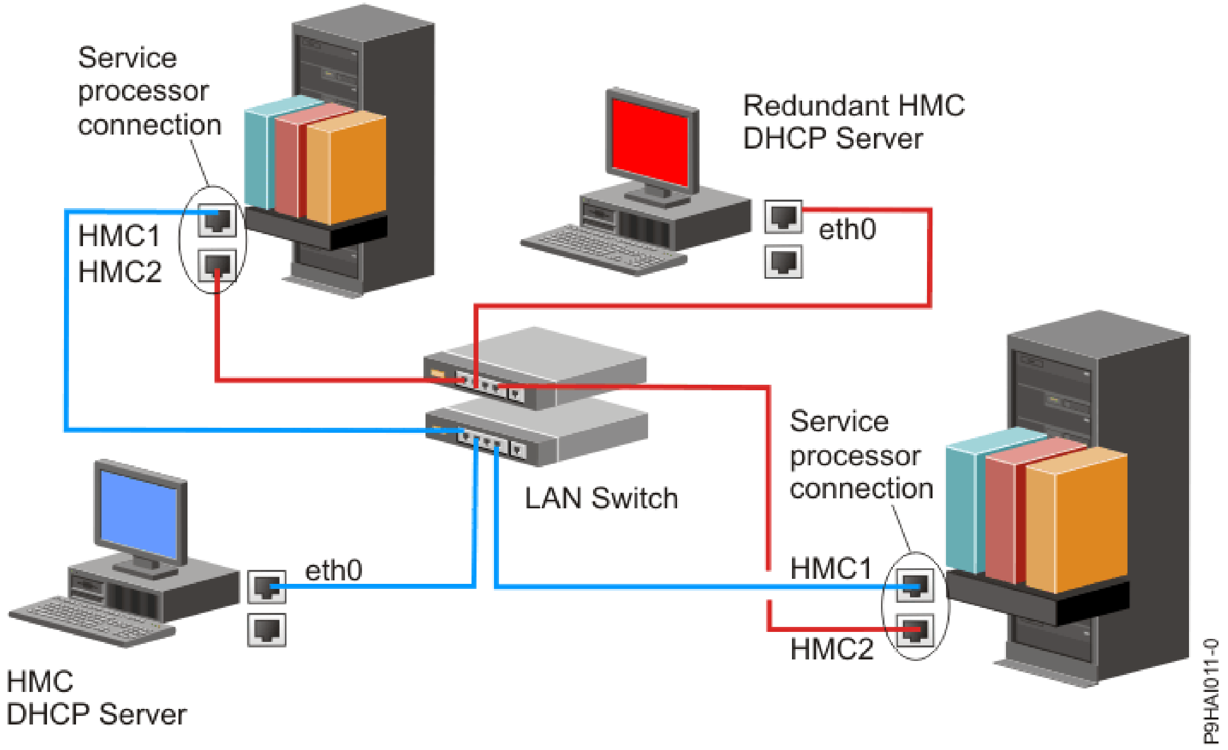
Service processor connection
HMC1
HMC2
Redundant HMC DHCP Server
eth0
LAN Switch
eth0
HMC DHCP Server
Service processor connection
HMC1
HMC2
P9HAI011-0

This figure shows a redundant HMC environment with two managed systems. The first HMC is connected to the first port on each FSP, and the redundant HMC is connected to the second port on each HMC. Each HMC is configured as a DHCP server, by using a different range of IP addresses. The connections are on separate private networks. As such, it is important to ensure that no FSP port is connected to more than one HMC.

Each managed system's FSP port that is connected to an HMC requires a unique IP address. To ensure that each FSP has a unique IP address, use the HMC's built-in DHCP server capability. When the FSP detects the active network link, it issues a broadcast request to locate a DHCP server. When correctly configured, the HMC responds to that request by allocating one of a selected range of addresses.

If you have multiple FSPs, you must have your own LAN switch or hub for the HMC to FSP private network. Alternately, this private segment can exist as several ports in a private *virtual LAN* (VLAN) on a larger managed switch. If you have multiple private VLANs, you must ensure that they are isolated and without any crossover traffic.

If you have more than one HMC, you must also connect each HMC to the logical partitions, and to each other, on the same open network.

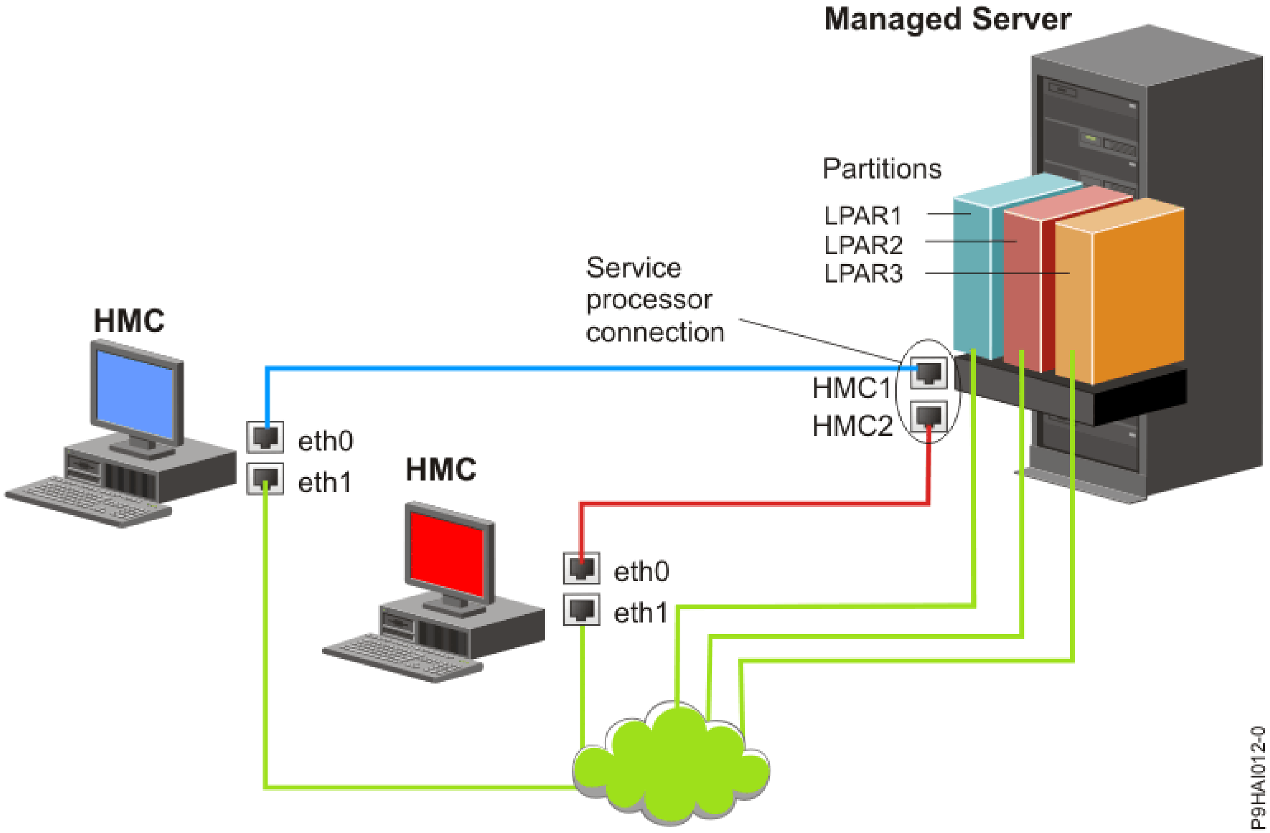

This figure shows two HMCs connected to a single managed server on the private network and to three logical partitions on the public network. You can have an extra Ethernet adapter for the HMC to have three network interfaces. You can use this third network as a management network or connect it to the CSM (Cluster Systems Manager) Management Server.

Deciding which connectivity method to use for the call-home server

Learn more about the connectivity options you have when you use the call-home server.

You can configure the Hardware Management Console (HMC) to send hardware service-related information to IBM by using a LAN-based internet connection, or a dial-up connection over a modem.

You have two communication choices when you configure the LAN-based internet connection. The first choice is to use standard Secure Sockets Layer (SSL). The SSL communication can be enabled to connect to the internet through your proxy server. SSL connectivity is more likely to be compliant with corporate security guidelines.

Note: If your open network interface connection uses only Internet Protocol Version 6 (IPv6), you cannot use internet VPN to connect to support. For more information about the protocols that are used, see "Choosing an Internet Protocol" on page 58.

The advantages to using an internet connection can include:

- Faster transmission speed
- Reduced customer expense (for example, the cost of a dedicated analog telephone line)
- Greater reliability

The following security characteristics are in effect, regardless of the connectivity method chosen:

- Remote Support Facility requests are always initiated from the HMC to IBM. An inbound connection is never initiated from the IBM Service Support System.
- All data that is transferred between the HMC and the IBM Service Support System are encrypted by using a high-grade encryption. Depending upon the connectivity method that is chosen, it is encrypted by using either SSL or IPSec Encapsulating Security Payload (ESP).

- When you initialize the encrypted connection, the HMC authenticates the target destination as the IBM Service Support System.

Data sent to the IBM Service Support System consists solely of information about hardware problems and configuration. No application or customer data is transmitted to IBM.

Using an indirect internet connection with a proxy server

If your installation requires the HMC to be on a private network, you might be able to connect indirectly to the internet by using an SSL proxy, which can forward requests to the internet. One of the other potential advantages of using an SSL proxy is that the proxy can support logging and audit facilities.

To forward SSL sockets, the proxy server must support the basic proxy header functions (as described in RFC 2616) and the CONNECT method. Optionally, basic proxy authentication (RFC 2617) can be configured so that the HMC authenticates before you attempt to forward sockets through the proxy server.

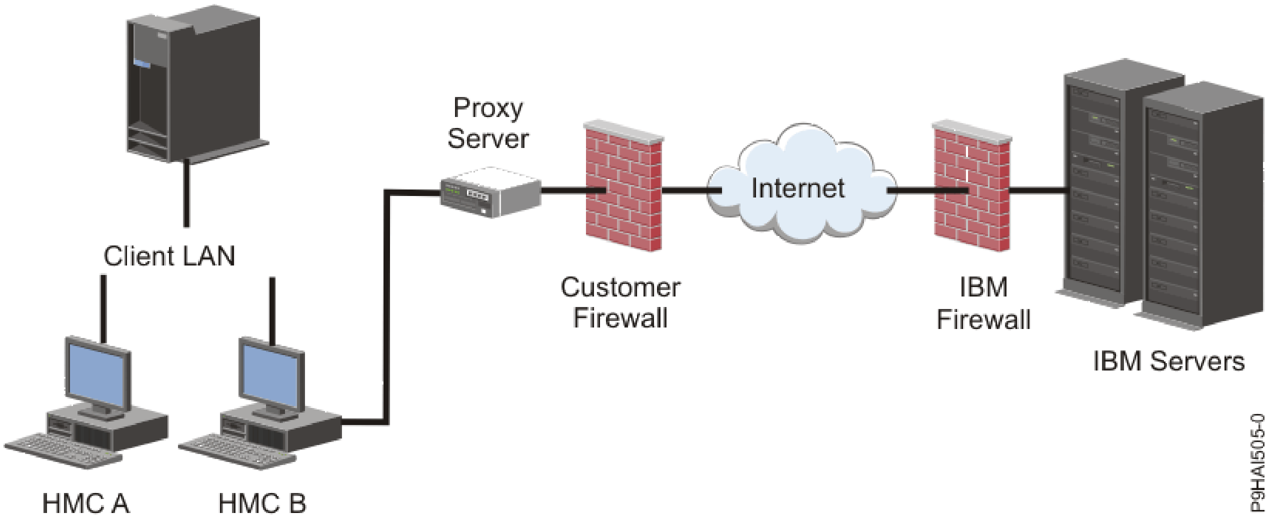

For the HMC to communicate successfully, the client's proxy server must allow connections to port 443. You can configure your proxy server to limit the specific IP addresses to which the HMC can connect. See "Internet SSL address lists" on page 58 for a list of IP addresses.

Using a direct internet SSL connection

If your HMC can be connected to the internet, and the external firewall can be set up to allow established TCP packets to flow outbound to the destinations described in "Internet SSL address lists" on page 58, you can use a direct internet connection.

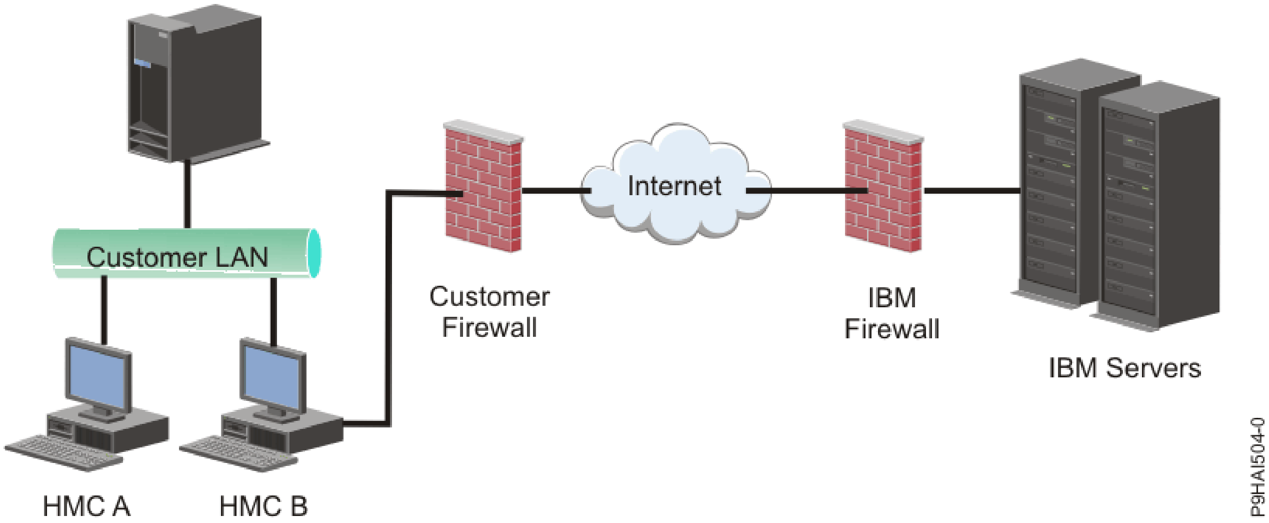

Using internet SSL to connect to remote support

All the communications are handled through TCP sockets that are initiated by the Hardware Management Console (HMC) and use a high-grade SSL to encrypt the data that is transmitted. The destination TCP/IP addresses are published (see "Internet SSL address lists" on page 58) so that external firewalls can be configured to allow these connections.

Note: The standard HTTPS port 443 is used for all communications.

The HMC can be enabled to connect directly to the internet or to connect indirectly from a proxy server that is provided by the customer. The decision about which approach is best for your installation depends on the security and networking requirements of your enterprise. The HMC (directly or through the SSL proxy) uses the following addresses when it is configured to use internet SSL connectivity.

Choosing an Internet Protocol

Determine the IP address version that is used when the Hardware Management Console (HMC) connects to your service provider.

Most users use Internet Protocol Version 4 (IPv4) to connect to a service provider. IPv4 addresses appear in the format that represents the 4 bytes of the IPv4 address, which is separated by periods (for example, 9.60.12.123) to access the internet. You can also use Internet Protocol Version 6 (IPv6) to connect to a service provider. IPv6 is often used by network administrators to ensure unique address space. If you are unsure of the Internet Protocol used by your installation, contact your network administrator. For more information about using each version, see "Setting the IPv4 address" on page 101 and "Setting the IPv6 address" on page 101.

Internet SSL address lists

Learn about the addresses that the Hardware Management Console (HMC) uses when the HMC is using internet SSL connectivity.

The HMC uses the following IPv4 addresses to contact IBM service and support when it is configured to use internet SSL connectivity.

The following IPv4 addresses are for all locations:

- 129.42.26.224
- 129.42.42.224
- 129.42.50.224
- 129.42.56.216
- 129.42.58.216
- 129.42.60.216
- 170.225.15.41

The following IPv4 addresses are for the Americas:

- 129.42.160.48
- 129.42.160.49
- 207.25.252.197
- 207.25.252.200
- 207.25.252.204

The following IPv4 addresses are for all locations other than the Americas:

- 129.42.160.48
- 129.42.160.50
- 207.25.252.197
- 207.25.252.200
- 207.25.252.205

Note: When you configure a firewall to allow an HMC to connect to these servers, only the IP addresses specific to the geographic region are required.

The HMC uses the following IPv6 addresses to contact IBM service and support when it is configured to use internet SSL connectivity:

- 2620:0:6C0:1::1000
- 2620:0:6C2:1::1000
- 2620:0:6C4:1::1000

Using multiple call-home servers

Learn about what you need to know when you decide to use more than one call-home server.

To avoid a single point of failure, configure the Hardware Management Console (HMC) to use multiple call-home servers. The first available call-home server attempts to handle each service event. If the connection or transmission fails with this call-home server, the service request is retried by using the other available call-home servers until one is successful or all servers are tried.

The connected HMC that is identified by the problem analysis to be the primary analyzing console for a given managed system that reports the problem. This primary console also replicates the problem report to any secondary HMC. This secondary HMC must be recognized on the network by the primary HMC. A secondary HMC is recognized by the primary HMC as an extra call-home server when:

- The primary HMC is configured to use "discovered" call-home servers and the call-home server is either on the same subnet as the primary HMC or it manages the same system.
- The call-home server is manually added to the list of call-home server consoles available for outbound connectivity.

Preparing for HMC configuration

Learn about the required configuration settings that you need to know before you begin the configuration steps.

To configure the HMC, you must understand the related concepts, make decisions, and prepare information.

Learn about the information that you need to connect your HMC to the following locations:

- Service processors in your managed systems
- Logical partitions on those managed systems
- Remote workstations
- IBM Service to implement “call-home” functions

To prepare for HMC configuration, complete the following steps:

1. Obtain and install the latest level of the HMC code version you want to install.
2. Determine the physical location of the HMC in relation to the servers it manages. If the HMC is more than 25 feet from its managed system, you must provide web browser access to the HMC from the managed system's location so that service personnel can access the HMC.
3. Identify the servers that the HMC manages.
4. Determine whether you use a private or an open network to manage servers. If you decide to use a private network, use DHCP, unless you are using a Cluster Systems Management (CSM) configuration. CSM does not support IPv6. To access CSM, you must have two networks. For more information about CSM, see the documentation that was provided with that feature. For more information about private and open networks, see “Selecting a private or open network” on page 74.
5. If you use an open network to manage an FSP, you must set the FSP's address manually through the Advanced System Management Interface menus. A private, non-routable network is recommended.
6. If you have two HMCs, designate a primary and secondary HMC. The primary HMC needs be physically closer to the system, and must be the HMC that is configured to call home.

7. Determine the network settings that you need to connect the HMC to remote workstations, logical partitions, and network devices.
8. Define how the HMC calls home. Call home options include either over an outbound-only Secure Socket Layer (SSL) internet connection, a modem, or a Virtual Private Network (VPN) connection.
9. Determine the HMC users that you create and their passwords, as well which roles they are given. You must assign the **hscroot** and **hscpe** users a password.
10. Document the following company contact information that is needed when you configure call home:
    - Company name
    - Administrator contact
    - Email address
    - Telephone numbers
    - Fax numbers
    - The street address of the HMC's physical location
11. If you plan to use email to notify operators or systems administrators when information is sent to IBM Service through call-home, identify the Simple Mail Transfer Protocol (SMTP) server and the email addresses you use.
12. You must define the following passwords:
    - The access password that is used to authenticate the HMC to the FSP.
    - The ASMI password that is used for the **admin** user.
    - The ASMI password that is used for the **general** user.

    Create the passwords when you connect from the HMC to a new server for the first time. If the HMC is a redundant or second HMC, obtain the HMC User password and be prepared to enter it when you connect the first time to the managed server's FSP.

When you complete these preparation steps, complete the "Preinstallation configuration worksheet for the HMC" on page 60.

Preinstallation configuration worksheet for the HMC

Use this worksheet to have the installation information you need ready for the installation.

Improved password policy for HMC

You must set a new password on the first use for newly manufactured systems with HMC version 9.940.0, or later, and after a factory reset of the system. This policy change helps to enforce that the HMC is not left in a state with a well-known password.

With HMC Version 9.940.0, and later, the `hscroot` password is expired and must be changed before you can access the functions of the HMC. For more information on how to change the password, see https://www.ibm.com/support/knowledgecenter/POWER9/p9eh6/p9eh6_useridsandpassword.htm. However, if you are upgrading from a previous HMC level or an operational installation, you do not have to change the password.

Network Settings

LAN Interface: Choose the available adapters (such as eth0, eth1) that is used by this HMC to connect to managed systems, logical partitions, service and support, and remote users. For more information, see "HMC network connections" on page 51. Connectivity from the HMC can either be on a private or open network.

Ethernet Adapter Speed and Duplex

Enter the wanted Ethernet adapter speed and duplex mode. The autodetection option determines which option is optimal if you are not sure which speed and duplex would produce optimum results for your hardware. Default = Autodetection Media speed specifies the speed in duplex mode of an

Ethernet adapter. Select Autodetection unless you need to specify a fixed media speed. Any device that is connected to the FSP (switches/HMC), must be set to Auto (Speed) / Auto (Duplex) mode, as it is the default FSP setting and cannot be changed.

Table 10. Ethernet Adapter Speed and Duplex

Characteristic	**eth0**	**eth1**	**eth2**	**eth3**
Select speed and duplex mode				
Media speed (Autodetection, 10/100/1000 Full/ Half Duplex)				

For more information about private and open networks, see "Private and open networks in the HMC environment" on page 53.

Table 11. Private or Open network

Characteristic	**eth0**	**eth1**	**eth2**	**eth3**
Specify **Private** or **Open** network for each adapter.				

Dynamic Host Configuration Protocol (DHCP) provides an automated method for dynamic client configuration. You can specify this HMC as a DHCP server. If this is the first or only HMC on the private network, enable the HMC as a DHCP server. When you enable the HMC as a DHCP server, the managed systems on the network are automatically configured and discovered by the HMC.

For Ethernet adapters specified as Private networks, complete the following table:

Table 12. DHCP server

Characteristics	**eth0**	**eth1**
Do you want to specify this HMC as a DHCP server? (yes/no)		
If yes, record the IP address range you want to use.		

If you are using the 7063-CR2 HMC, you must connect the Ethernet **IPMI** port to a network to access the baseboard management controller (BMC) on the HMC. For more information, see "Configure BMC connectivity (7063-CR2)" on page 100. Complete the following table for your BMC connection.

Table 13. BMC connection

Characteristics	**IPMI**
Do you want to configure this connection through DHCP mode? (yes/no)	
If no, list the specified static addresses below:	
IP address:	
Subnet mask:	
Gateway:	

If you are using the 7063-CR1 HMC, you must connect the Ethernet **IPMI** port to a network to access the baseboard management controller (BMC) on the HMC. For more information, see "Configure BMC connectivity (7063-CR1)" on page 100. Complete the following table for your BMC connection.

Table 14. BMC connection

Characteristics	IPMI
Do you want to configure this connection through DHCP mode? (yes/no)	
If no, list the specified static addresses below:	
IP address:	
Subnet mask:	
Gateway:	

For Ethernet adapters specified as *open* networks, complete the following tables. For more information about the different Internet Protocol versions, see "Configuring the HMC network types" on page 70.

Using IPv6

If you are using IPv6, talk to your network administrator and decide how you want to obtain IP addresses. Then, complete the following tables:

Table 15. IPv6 (static)

Characteristic	eth0	eth1	eth2	eth3
Are you using a statically assigned IP address? If yes, record that address here.				

Table 16. IPv6 (DHCP server)

Characteristic	eth0	eth1	eth2	eth3
Are you getting IP addresses from a DHCP server? (Yes/No)				

Table 17. IPv6 (IPv6 router)

Characteristic	eth0	eth1	eth2	eth3
Are you getting IP addresses from an IPv6 router?				

For more information about setting IPv6 addresses, see "Setting the IPv6 address" on page 101. For more information about using only IPv6 addresses, see "Using only IPv6 addresses" on page 101.

Using IPv4

Complete the following tables for Ethernet adapters that are specified as open networks by using IPv4.

Table 18. IPv4

Characteristics	eth0	eth1	eth2	eth3
Do you want to obtain an IP address automatically? (yes/no)				
If no, list the specified address below:				
TCP/IP Interface Address:				
TCP/IP Interface Network Mask:				
Firewall Settings:				
Would you like to configure HMC firewall settings? (yes/no)				
If yes, list the applications and IP addresses that must be allowed through the firewall:				

TCP/IP information
A unique TCP/IP address is required for each node, both for Support Element (SE) and Hardware Management Console (HMC). The assigned network mask is used to generate a unique address, by default, for the local private LAN. If the nodes are connected into a larger network with an administered TCP/IP address, you can specify the TCP/IP address to be used. The default is generated by the system.

Firewall settings
HMC firewall settings create a security barrier that allows or denies access to specific network applications on the HMC. You can specify these control settings individually for each physical network interface, enabling you control over which HMC network applications can be accessed on each network.

If you configure at least one adapter as an Open network adapter, you must provide the following additional information to enable your HMC to access the LAN:

Table 19. Open network adapter

Local host information	
HMC host name:	
Domain name:	
Description of HMC:	
Gateway information	
Gateway Address: (nnn.nnn.nnn.nnn)	
Gateway device:	
DNS enablement	

Table 19. Open network adapter (continued)

Local host information	
Do you want to use DNS? (yes/no)	
If “yes”, specify DNS Server Search Order below:	
1.	
2.	
Domain suffix search order:	
1.	
2.	

Local Host information
To identify your Hardware Management Console (HMC) to the network, enter the HMC's host name and domain name. Unless you are using only short host names on your network, enter a fully qualified host name. Domain name example: name.yourcompany.com

Gateway information
To define a default gateway, complete the TCP/IP address to be used for routing IP packets. The gateway address informs each computer or network device when to send data if the target station is not on the same subnet as the source.

DNS Enablement
The Domain Name System (DNS) is used to provide a standard naming convention for locating IP-based computers. By defining DNS servers, you can use host names to identify servers and Hardware Management Consoles (HMCs) rather than IP addresses.

DNS Server Search Order
Enter the IP addresses of DNS servers to be searched for mapping the host names and IP addresses. This search order is available only when DNS is enabled.

Domain Suffix Search Order
Enter the domain suffixes you are using. The HMC uses domain suffixes to append to unqualified names for DNS searches. Suffixes are searched in the order in which they are listed. This search order is available only when DNS is enabled.

Email notification

List email contact information if you want to be notified by email when hardware problem events occur on your system.

Table 20. Email notification

Characteristics	**Entry field**
Email Addresses:	
SMTP server:	
Port:	
Errors to be notified:	
Only call-home problem events	
All problem events	

SMTP server
Type the simple mail transfer Protocol (SMTP) address of the server to be notified of a system event. An example of an SMTP server name is `relay.us.ibm.com`.

SMTP is the Protocol that is used to send email. When you use SMTP, a client sends a message and communicates with the SMTP server by using the SMTP Protocol.

If you do not know the SMTP address of your server or are not sure, contact your network administrator.

Port
Type the port number of the server to be notified of a system event, or use the default port.

Email addresses to be notified
Enter configured email addresses to be notified when a system event occurs.

- Select **Only call-home problem events** to receive notification only when events occur that create a call-home function.
- Select **All problem events** to receive notification when any events occur.

Service contact information

Table 21. Service contact information

Characteristics	Entry field
Company name	
Administrator name	
Email address	
Phone number	
Alternative phone number	
Fax number	
Alternative phone number	
Street address	
Street address 2	
City or locality	
State	
Postal code	
Country or region	
Location of HMC (if same as above administrator address, specify "same"):	
Street address	
Street address 2	
City or locality	
State	
Postal code	
Country or region	

Service authorization and connectivity

Select the type of connection to contact your service provider. For a description of these methods that include security characteristics and configuration requirements, see "Choosing existing call-home servers to connect to service and support for this HMC" on page 107.

Table 22. Service authorization and connectivity

Characteristics	Entry field
Secure Sockets Layer (SSL) through the internet	_____
Virtual private network (VPN) through the internet	_____

Secure Sockets Layer (SSL) through the internet:

If you have an existing internet connection from your HMC, you can use it to call your service provider. You can connect directly to your service provider by using encrypted Secure Sockets Layer (SSL) by using the existing internet connection. Select **Use SSL Proxy** if you want to configure the use of encrypted SSL by using an indirect connection that uses an SSL Proxy.

Table 23. SSL

Characteristics	Entry field
Use SSL proxy? (yes/no)	
If yes, list information below:	
Address:	
Port:	
Authenticate with the SSL Proxy?	
If yes, list information below:	
User:	
Password:	

Internet connection Protocol used

For more information about the different internet Protocols, see "Configuring the HMC network types" on page 70.

___ IPv4

___ IPv6

___ IPv4 and IPv6

Virtual Private Network (VPN)

If you have an existing internet connection from your HMC, you can use it to call your service provider. You can connect directly to your service provider by virtual private network (VPN) by using the existing internet connection.

Note: If you select Virtual Private Network (VPN) through the internet, you cannot select any other options.

Call-home servers

Determine which HMCs you want to configure to connect to service and support as call-home servers. For more information about using multiple call-home servers, see "Using multiple call-home servers" on page 59.

___ This HMC

___ Another HMC

If you checked **Another HMC**, list the other HMCs that are configured as call-home servers here:

Table 24. Other HMCs that are configured as call-home servers

List of HMC host names or IP addresses that are configured as call-home servers

Extra Support Benefits

My Systems and Premium Search

Table 25. My Systems and Premium Search

Characteristics	**Entry field**
List your IBM ID	__________________________
List any additional IBM IDs	__________________________

To access valuable, customized support information in the My Systems and Premium Search sections of the Electronic Services website, Customers must register their IBM ID with this system. If you do not already have one, you can register for an IBM ID at: www.ibm.com/account/profile.

Note: IBM provides personalized web functions that use information that is collected by the IBM Electronic Service Agent application. To use these functions, you must first register on the IBM Registration website at http://www.ibm.com/account/profile.

To authorize users to use the Electronic Service Agent information to personalize the web functions, enter your IBM ID that you registered on the IBM Registration website. Go to http://www.ibm.com/support/electronic to see the valuable support information available to customers that register an IBM ID with their systems.

Configuring the HMC

Learn how to configure network connections, security, service applications, and some user preferences.

Depending on the level of customization you intend to apply to your HMC configuration, you have several options for setting up your HMC to suit your needs. The Guided Setup wizard is a tool on the HMC designed to ease the setup of the HMC. You can choose a fast path through the wizard to quickly create the recommended HMC environment, or you can choose to fully explore the available settings that the wizard guides you through. You can also perform the configuration steps without the aid of the wizard by Configuring the HMC by using the HMC menus.

Before you start, gather the required configuration information that you need to complete the steps successfully. See "Preparing for HMC configuration" on page 59 for a list of the required information. When you are finished preparing, ensure that you complete the "Preinstallation configuration worksheet for the HMC" on page 60 and then return to this section.

Configuring the HMC by using the menus

This section provides a complete list of all HMC configuration tasks, guiding you through the process of configuring your HMC. Choose this option if you prefer not to use the Guided Setup wizard.

You must restart your HMC for the configuration settings to take effect, so you might want to print this checklist and keep it with you as you configure your HMC.

This information contains references to tasks that are not included in this document. You can access the IBM Power Systems hardware information on the HMC or on the Web. On the HMC, IBM Knowledge Center can be accessed from the upper-right corner of the task bar. On the web, IBM Knowledge Center can be accessed at https://www.ibm.com/support/knowledgecenter.

This information contains references to tasks that are not included in this PDF. You can access additional support materials by referring to the **Additional Resources** section on the HMC Welcome page.

Prerequisites

Before you begin configuring the HMC using the HMC menus, be sure to complete the configuration preparation activity described in "Preparing for HMC configuration" on page 59.

Table 26. Manual HMC configuration tasks and where to find related information

Task	Where to find related information
1. Start the HMC.	"Starting the HMC" on page 69
2. Set the date and time.	
3. Change predefined passwords.	
4. Create additional users and return to this checklist when you have completed this step.	
5. Configure network connections.	"Configuring the HMC network types" on page 70
6. For HMC model 7063-CR1, you must configure the baseboard management controller (BMC) IP address.	"Configure BMC connectivity (7063-CR1)" on page 100
7. If you are using an open network and a fixed IP address, set identification information.	
8. If you are using an open network and a fixed IP address, configure a routing entry as the default gateway.	"Configuring a routing entry as the default gateway" on page 103
9. If you are using an open network and a fixed IP address, configure domain name services.	"Configuring domain name services" on page 103
10. If you are using a fixed IP address and have DNS enabled, configure domain suffixes.	"Configuring domain suffixes" on page 104
11. Configure your server to connect to IBM service and support and return to this checklist when you have completed this step.	"Configuring the local console to report errors to service and support" on page 106
12. Configure the Events Manager for Call Home.	"Configuring the Events Manager for Call Home" on page 109
13. Connect the managed system to a power source.	
14. Set passwords for the managed system, and each of the ASMI passwords (general and admin)	"Setting passwords for the managed system" on page 110
15. Access ASMI to set the date and time on the managed system.	
16. Start the managed system and return to this checklist when you have completed this step.	
17. Ensure that you have one logical partition on the managed system.	
18. Optional: add another managed system and return to this checklist when you have completed this step.	

Table 26. Manual HMC configuration tasks and where to find related information (continued)

Task	Where to find related information
19. Optional: If you are installing a new server with your HMC, configure the logical partitions and install the operating system.	
20. If you are not installing a new server at this time, perform optional postconfiguration tasks to further customize your configuration.	"Postconfiguration steps" on page 111

Starting the HMC

You can long in to the HMC and choose which language you want to be displayed in the interface. Use the default User ID `hscroot` and password `abc123` to log on to the HMC for the first time.

About this task

To start the HMC, do the following procedure:

Procedure

1. Turn on the HMC by pressing the power button.
2. If English is your language preference, continue with step 4.

   If your language preference is a language other than English, type the number **2** when you are prompted to change the locale.

   Note: This prompt times out in 30 seconds if you do not act.
3. Select the locale that you want to display from the list in the **Locale Selection** window, and click **OK**. The locale identifies the language that the HMC interface uses.
4. Click **Log on and launch the Hardware Management Console web application**.
5. Log in to the HMC with the following default user ID and password:

   ID: `hscroot`
   Password: `abc123`

   HMC Enhanced
   Displays the newer enhanced GUI with the enhanced PowerVM features.

   HMC Classic
   Displays the standard GUI without the enhanced PowerVM features.

   Note: When the HMC is working as a DHCP server, the HMC uses the default password when it connects to the service processor for the first time.
6. Press Enter.

Changing the date and time

The battery-operated clock keeps the date and time for the Hardware Management Console (HMC). You might need to reset the console date and time if the battery is replaced, or if you physically move your system to a different time zone. Learn how to change the date and time for the HMC.

About this task

If you change the date and time information, the change does not affect the systems and logical partitions that the HMC manages.

To change the date and time for the HMC, complete the following steps:

Procedure

1. Ensure that you are a member of one of the following roles:
   - Super administrator
   - Service representative
   - Operator
   - Viewer
2. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
3. In the content pane, click **Change Date and Time**.
4. If you select **UTC** in the **Clock** field, the time setting adjusts automatically for Daylight Saving Time in the time zone you select. Enter the date, time, and time zone, and click **OK**.

Results

Configuring the HMC network types

Configure your HMC so that it can communicate to the managed system, logical partitions, remote users, and service and support.

Configuring HMC settings to use an open network to connect to the managed system

Configure the HMC so that it can connect to and manage a managed system using an open network.

Before you begin

To configure the HMC network settings so that it can connect to the managed system using an open network, do the following:

Table 27. Configuring HMC settings to use an open network to connect to the managed system

Task	**Where to find related information**
1. Decide which interface you want to use for your managed system. **eth0** is preferred.	“Preinstallation configuration worksheet for the HMC” on page 60
2. Identify the Ethernet ports for your HMC.	“Identifying the Ethernet port that is defined as eth0” on page 73
3. Configure the Ethernet adapter by performing the following tasks:	
a. Set the media speed.	“Setting the media speed” on page 74
b. Select the open network type.	“Selecting a private or open network” on page 74
c. Set static addresses.	“Setting the IPv6 address” on page 101
d. Set the firewall.	“Changing HMC firewall settings” on page 102
e. Configure the default gateway.	“Configuring a routing entry as the default gateway” on page 103
f. Configure DNS.	“Configuring domain name services” on page 103
4. Configure additional adapters, if you have them.	
5. Test the connection between the managed server and the HMC.	“Testing the connection between the HMC and the managed system” on page 111

Configuring HMC settings to use a private network to connect to the managed system
Configure the HMC so that it can connect to and manage a managed system using a private network.

Before you begin

To configure the HMC network settings so that it can connect to the managed system using a private network, do the following:

Table 28. Configuring HMC settings to use a private network to connect to the managed system

Task	Where to find related information
1. Decide which interface you want to use for your managed system.	"Preinstallation configuration worksheet for the HMC" on page 60
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 73
3. Configure the HMC as a DHCP server.	"Configuring the HMC as a DHCP server" on page 75
4. Test the connection between the managed server and the HMC.	"Testing the connection between the HMC and the managed system" on page 111

Configuring HMC settings to use an open network to connect to logical partitions

Before you begin

To configure the HMC network settings so that it can connect to logical partitions using an open network, do the following:

Table 29. Configuring HMC settings to use an open network to connect to logical partitions

Task	Where to find related information
1. Decide which interface you want to use for your managed system.	"Preinstallation configuration worksheet for the HMC" on page 60
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 73
3. Configure the Ethernet adapter by performing the following tasks:	
a. Set the media speed.	"Setting the media speed" on page 74
b. Select the open network type.	"Selecting a private or open network" on page 74
c. Set static addresses.	"Setting the IPv6 address" on page 101
d. Set the firewall.	"Changing HMC firewall settings" on page 102
e. Configure the default gateway.	"Configuring a routing entry as the default gateway" on page 103
f. Configure DNS.	"Configuring domain name services" on page 103
4. Configure additional adapters, if you have them.	
5. Test the connection between the managed server and the HMC.	"Testing the connection between the HMC and the managed system" on page 111

Configuring HMC settings to use an open network to connect to remote users

Before you begin

To configure the HMC network settings so that it can connect to remote users using an open network, do the following:

Table 30. Configuring HMC settings to use an open network to connect to remote users

Task	Where to find related information
1. Decide which interface you want to use for your managed system.	"Preinstallation configuration worksheet for the HMC" on page 60
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 73
3. Configure the Ethernet adapter by performing the following tasks:	
a. Set the media speed.	"Setting the media speed" on page 74
b. Select the open network type.	"Selecting a private or open network" on page 74
c. Set static addresses.	"Setting the IPv6 address" on page 101
d. Set the firewall.	"Changing HMC firewall settings" on page 102
e. Configure the default gateway.	"Configuring a routing entry as the default gateway" on page 103
f. Configure DNS.	"Configuring domain name services" on page 103
g. Configure suffixes.	"Configuring domain suffixes" on page 104
4. Configure additional adapters, if you have them.	

Configuring HMC call-home server settings

Before you begin

To configure the HMC call-home server settings so that problems can be reported, do the following:

Table 31. Configuring HMC call-home server settings

Task	Where to find related information
1. Be sure you have all the required customer information	"Preinstallation configuration worksheet for the HMC" on page 60
2. Configure this HMC to report errors or choose an existing call-home server to report errors	"Configuring the local console to report errors to service and support" on page 106 "Choosing existing call-home servers to connect to service and support for this HMC" on page 107
3. Verify that your call-home configuration is working	"Verifying that your connection to service and support is working" on page 108
4. Authorize users to view collected system data	"Authorizing users to view collected system data" on page 108
5. Schedule transmission of system data	"Transmitting service information" on page 109

Identifying the Ethernet port that is defined as eth0
Your Ethernet connection to the managed server must be made by using the Ethernet port that is defined as `eth0` on your HMC.

If you did not install any additional Ethernet adapters in the PCI slots on your HMC, then the primary-integrated Ethernet port is always defined as `eth0` or `eth1` on your HMC, if you intend to use the HMC as a DHCP server for your managed systems.

If you install extra Ethernet adapters in the PCI slots, then the port that is defined as `eth0` depends on the location and type of Ethernet adapters that are installed.

Note: The following general rules might not apply for all configurations.

The following table describes the rules for Ethernet placement by HMC type.

Table 32. HMC types and associated rules for Ethernet placement

HMC type	**Rules for Ethernet placement**
Rack-mounted HMCs with two integrated Ethernet ports.	The HMC supports only one extra Ethernet adapter. • If an extra Ethernet adapter is installed, then that port is defined as `eth0`. In this case, the primary-integrated Ethernet port is then defined as `eth1`, and the secondary integrated Ethernet port is defined as `eth2`. • If the Ethernet adapter is a dual port Ethernet adapter, then the port that is labeled Act/Link A is `eth0`. The port that is labeled `Act/link B` is `eth1`. In this case, the primary-integrated Ethernet port is then defined as `eth2`, and the secondary integrated Ethernet port is defined as `eth3`. • If no adapters are installed, then the primary-integrated Ethernet port is defined as `eth0`.
Stand-alone models with a single integrated Ethernet port.	The definitions depend upon the type of Ethernet adapter that is installed: • If only one Ethernet adapter is installed, then that adapter is defined as `eth0`. • If the Ethernet adapter is a dual port Ethernet adapter, then the port that is labeled `Act/link A` is `eth0`. The port that is labeled `Act/link B` would be `eth1`. In this case, the primary-integrated Ethernet port is then defined as `eth2`. • If no adapters are installed, then the integrated Ethernet port is defined as `eth0`. • If multiple Ethernet adapters are installed, see "Determining the interface name for an Ethernet adapter" on page 73.

Determining the interface name for an Ethernet adapter
If you configure the HMC as a DHCP server, that server can operate only on the network interface card (NIC) connectors that the HMC identifies as `eth0` and `eth1`. You might also need to determine which NIC

connector you need to plug the Ethernet cable into. Learn more about determining which NIC connectors the HMC identifies as `eth0` and `eth1`.

About this task

To determine the name the HMC has assigned to an Ethernet adapter, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Settings**.
2. In the content pane, click **Change Network Settings**.
3. From the **Change Network Settings** window, click the **LAN adapters** tab. The following example entry shows that this Ethernet port is identified as `eth0`: `Ethernet eth0 52:54:00:fa:b6:8e (<IP address of HMC>)`.
4. Record your results. If you need to view or change the LAN adapter settings, click **Details**.
5. Click **OK**.

Setting the media speed
Learn how to specify the media speed that includes the speed and duplex mode of the Ethernet adapter.

Before you begin

The default for the HMC adapter settings is **Autodetection**. If this adapter is connected to a LAN switch, you must match the switch port settings. To set the media speed and duplex, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. In the local area network (LAN) information section, select **Autodetection** or the appropriate media speed and duplex combination.
6. Click **OK**.

Selecting a private or open network
A *private service network* consists of the Hardware Management Console (HMC) and the managed systems. A private service network is restricted to consoles and the systems they manage, and is separate from your company network. An *open network* consists of your private service network and your company network. An open network might contain network endpoints in addition to consoles and managed systems, and might span across multiple subnets and network devices.

About this task

To select a private or public network, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Settings**.

2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **LAN Adapter** tab.
6. In the local area network information page, select **Private** or **Open**.
7. Click **OK**.

Configuring the HMC as a DHCP server
Dynamic Host Configuration Protocol (DHCP) provides an automated method for dynamic client configuration.

To configure the Hardware Management Console (HMC) as a DHCP server, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Customize Network Settings window opens.
3. Select the LAN adapter that you want to work with and click **Details**.
4. Select **Private** and then select the network type.
5. In the DHCP Server section, select **Enable DHCP Server** to enable the HMC as a DHCP server.

   Note: You can configure the HMC to be a DHCP server only on a private network. If you use an open network, the option to select the **Enable DHCP** is not available.
6. Enter the address range of the DHCP server.
7. Click **OK**.

If you configured your HMC to be a DHCP server on a private network, you must verify that your HMC DHCP private network is configured correctly. For information about connecting your HMC to a private network, see “Selecting a private or open network” on page 74.

For more information, see “ HMC as a DHCP server” on page 54.

Managing the system by using OpenBMC-based HMC (7063-CR2)
IBM Power Systems servers use a baseboard management controller (BMC) for system service management, monitoring, maintenance, and control. The BMC also provides access to the system event log files (SEL). The BMC is a specialized service processor that monitors the physical state of the system by using sensors. A system administrator or service representative can communicate with the BMC through an independent connection. The OpenBMC tool provides a communication method to the BMC, by using a command-line interface. The OpenBMC tool can be used either from a remote Linux system, or from the host operating system console window. The OpenBMC tool can be connected remotely to the BMC by using a configured Ethernet port. You can connect your server to a monitor by using the VGA port at the rear of the server.

Managing the system by using the OpenBMC tool
Learn how to configure and manage your system by using the OpenBMC tool.

Downloading and installing the OpenBMC tool
Learn how to download and install the OpenBMC tool.

About this task

To download and install the OpenBMC tool, complete the following steps:

Procedure

1. Go to the IBM Support Portal.
2. In the search field, type: `Scale-out LC System Event Log Collection Tool`.

3. Click the **Scale-out LC System Event Log Collection Tool** entry and follow the instructions to install and run the OpenBMC tool.

Basic commands and functionality of the OpenBMC tool
The OpenBMC tool provides support for working with system event logs, updating system firmware, identifying the system, powering off the system, and other service-related functions.

OpenBMC tool top-level options
Learn more about the top-level options for the OpenBMC tool commands.

About this task

- `-H`: Host name or IP address of the BMC.
- `-U`: User name to log in with.
- `-A`: Provides a prompt to ask for the password.
- `-P`: Password for the user name.
- `-j`: Change output format to JSON.
- `-t`: Location of the policy table to use.
- `-T`: Provides time statistics for logging in, running the command, and logging out.
- `-V`: Displays current version of the OpenBMC tool.

System event log commands
Learn more about system event log commands for the OpenBMC tool.

Procedure

- To print a list of the system event logs in a readable format, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  sel print
  ```
- To list the system event logs in raw data, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  sel list
  ```
- To change the status of a system event log to resolved, use the following command:

  `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel resolve -n x`, where *x* is the system event log number.
- To collect all service data including system event logs, use the following command:

  `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> collect_service_data`.
- To clear gard records for disable hardware, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  gardclear
  ```
- To clear the alert logs of entries, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  sel clear
  ```

System firmware update command
Learn more about the system firmware update command.

Procedure

- To update the system firmware, use the following command:

`openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> firmware flash <bmc or pnor> -f xxx.tar`, where *bmc* or *pnor* is the type of image you wish to flash to the system.

Note: If you are not in the same folder as the TAR file, you must include the full path to the folder where the file resides.

- To activate a firmware image that is available in the BMC, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
firmware activate <firmware image ID>
```

System identify commands
Learn more about the system identify commands.

Procedure

- To activate the blue system identify LED, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis identify on
```

- To turn off the blue system identify LED, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis identify off
```

- To check the status of the blue system identify LED, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis identify status
```

System power on and power off commands
Learn more about the system power on and power off commands.

Procedure

- To check the power status of the system, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis power status
```

- To power on the system, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis power on
```

- To power off the system normally, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis power softoff
```

- To power off the system immediately, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis power hardoff
```

System sensor commands
Learn more about the system sensor commands.

Procedure

- To display a list of all monitoring sensors, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
sensors print
```

or

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
sensors list
```

System FRU commands
Learn more about the system FRU commands.

Procedure

- To display a list of all inventory items, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  fru print
  ```

  or

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  fru list
  ```
- To display the known status of all FRU items, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  fru status
  ```

  Note: The FRU item must be designated as a replaceable FRU by the BMC.
- To automate the review of FRU status commands and to determine if there is a performance impact on the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  health_check
  ```

  Note: This command does not guarantee a healthy system as there can be system event logs entries that are not associated with the inventory items.

System BMC reset commands
Learn more about the system BMC reset commands.

Procedure

- To do a warm reset of the BMC remotely and without an AC cycle, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  bmc reset warm
  ```
- To do a cold reset of the BMC remotely and without an AC cycle, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  bmc reset cold
  ```

System dump commands
Learn more about the system dump commands.

Procedure

- To create a new dump file, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump create
  ```
- To list all dump files in the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump list
  ```
- To delete a specific dump file from the system, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
dump delete -n <dump file entry>
```

- To delete all dump files from the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump delete all
  ```

- To retrieve a specific dump file, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump retrieve -n <dump file entry>
  ```

- To retrieve a dump file and save it to specific directory, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump retrieve -s <location to save dump file>
  ```

  Note: If you do not specify a location, the file is saved in the OS where the command is run in the temp directory.

Enabling and disabling local BMC user accounts
Learn more about the **local_users**commands.

About this task

The local user accounts on the BMC, such as root, can be disabled, queried, and re-enabled with the **local_users** sub-command.

Note: After disabling local users, the LDAP user needs to be available for further interaction with the BMC, including enabling local users by using OpenBMC tool.

Procedure

- To view current local user account status, use the following command:

  ```
  openbmctool <connection options> local_users queryenabled
  ```

- To disable all local user accounts, use the following command:

  ```
  openbmctool <connection options> local_users disableall
  ```

- To enable all local user accounts, use the following command:

  ```
  openbmctool <connection options> local_users enableall
  ```

Remote logging by using rsyslog
Learn more about the remote logging commands.

About this task

The BMC can stream out local logs (that go to the systemd journal) by using RSYSLOG. The BMC sends everything in the logs. Any kind of filtering and appropriate storage has to be managed on the rsyslog server.

Procedure

- To configure the rsyslog server for remote logging, use the following command:

  ```
  openbmctool <connection options> logging remote_logging_config -a <IP
  address> -p <port>
  ```

  Note: The IP address and port are for the remote rsyslog server. After the command is run, the remote rsyslog server starts to receive logs from the BMC.

- To disable remote logging, use the following command:

  ```
  openbmctool <connection options> logging remote_logging disable
  ```

Note: Disable remote logging before you switch remote logging from an existing remote server to a new one.

- To view the remote logging configuration, use the following command:

  ```
  openbmctool <connection options> logging remote_logging view
  ```

  Note: This command prints out the IP address and port of the remote rsyslog server in JavaScript Object Notation (JSON) format.
- To turn REST API logging on, use the following command:

  ```
  openbmctool <connection options> logging rest_api on
  ```
- To turn REST API logging off, use the following command:

  ```
  openbmctool <connection options> logging rest_api off
  ```

  Note: REST API logging is turned off by default.

Certificate management
Learn more about the certificate management commands.

About this task

You can replace the existing certificate and private key file with another (possibly CA signed) certificate and private key file. You can install server, client, and root certificates.

Procedure

- To update the HTTPS server certificate, use the following command:

  ```
  openbmctool <connection options> certificate update server https -f <File>
  ```

  Note: The `<File>` is the privacy-enhanced mail (PEM) file that contains both the certificate and the private key.
- To update the LDAP client certificate, use the following command:

  ```
  openbmctool <connection options> certificate update client ldap -f <File>
  ```

  Note: The `<File>` is the PEM file that contains both the certificate and the private key.
- To update the LDAP root certificate, use the following command:

  ```
  openbmctool <connection options> certificate update authority ldap -f <File>
  ```

  Note: The `<File>` is the PEM file that contains only the certificate.
- To delete the HTTPS server certificate, use the following command:

  ```
  openbmctool <connection options> certificate delete server https
  ```

  Note: Deleting a certificate creates and installs a new self-signed certificate.
- To delete the LDAP client certificate, use the following command:

  ```
  openbmctool <connection options> certificate delete client ldap
  ```
- To delete the LDAP root certificate, use the following command:

  ```
  openbmctool <connection options> certificate delete authority ldap
  ```

  Note: Deleting the root certificate can cause an LDAP service outage.

LDAP configuration
Learn more about the LDAP configuration commands.

About this task

In the BMC, LDAP is used for remote authentication. The BMC does not support remote user-management functionality. The BMC supports both secure and non-secure LDAP configuration.

Procedure

- To create the LDAP configuration (non-secure), use the following command:

  ```
  openbmctool.py <connection options> ldap enable --uri="ldap://
  <ldap server IP/hostname>" --bindDN=<bindDN> --baseDN=<basDN> --
  bindPassword=<bindPassword> --scope="sub/one/base" --serverType="OpenLDAP/
  ActiveDirectory"
  ```

  Note: Configuring a fully qualified domain name or hostname in the `uri` parameter requires the domain name system (DNS) server to be configured on the BMC.

- To create the LDAP configuration (secure), use the following command:

  ```
  openbmctool.py <connection options> ldap enable --uri="ldaps://
  <ldap server IP/hostname>" --bindDN=<bindDN> --baseDN=<basDN> --
  bindPassword=<bindPassword> --scope="sub/one/base" --serverType="OpenLDAP/
  ActiveDirectory"
  ```

  Notes:

  1. It is common to encounter the following error when you run the above `openbmctool.py` command string:

     xyz.openbmc_project.Common.Error.NoCACertificate

     This error means that the BMC client needs to verify that the LDAP server certificate is signed by a known certification authority (CA). An administrator needs to upload the CA certificate to the BMC to resolve this error.

  2. The OpenBMC tool does not support individual LDAP configuration property updates. To update a single property, the administrator must recreate the LDAP configuration with the changed values.

- To delete the LDAP configuration, use the following command:

  ```
  openbmctool.py <connection options> ldap disable
  ```

  Note: The root user must be enabled before you run the command, otherwise the BMC is not accessible. To enable all local user accounts, see Enabling and disabling local user accounts.

- To add privilege mapping use the following command:

  ```
  openbmctool.py <connection options> ldap privilege-mapper create --
  groupName=<groupName> --privilege="priv-admin/priv-user"
  ```

- To delete privilege mapping, use the following command:

  ```
  openbmctool.py <connection options> ldap privilege-mapper delete --
  groupName=<groupName>
  ```

- To list privilege mapping, use the following command:

  ```
  openbmctool.py <connection options> ldap privilege-mapper list
  ```

  The normal workflow for LDAP configuration is in the following order:

  1. Configure the DNS server.
  2. Configure LDAP.
     a. Configure the CA certificate for secure LDAP configuration.

   b. Create LDAP configuration with local user.

3. Configure user privilege.

Notes:

1. If you login with LDAP credentials and have not added privilege mapping for the LDAP credentials, then you will get the following error message:

   403, 'LDAP group privilege mapping does not exist'.

   You can avoid this error by adding privilege mapping.

2. The following error message might mean that user lacks sufficient privileges on the BMC:

   Insufficient privileges

   You can avoid this error by adding privilege mapping.

3. After you setup the LDAP, the OpenBMC tool connection options work with both LDAP and local users.

Network configuration
Learn more about the network configuration commands.

Procedure

- To enable DHCP, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network enableDHCP -I
  <Interface name>
  ```

- To disable DHCP, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network disableDHCP -I
  <Interface name>
  ```

- To get the host name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getHostName
  ```

- To set the host name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setHostName -H
  <host name>
  ```

- To get the domain name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getDomainName
  -I <Interface name>
  ```

- To set the domain name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setDomainName
  -I <Interface name> -D DomainName1,DomainName2,..
  ```

- To get the media access control (MAC) address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getMACAddress
  -I <Interface name>
  ```

- To set the MAC address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setMACAddress
  -I <Interface name> -MA xx:xx:xx:xx:xx
  ```

- To get the default gateway, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getDefaultGW
  ```

- To set the default gateway, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setDefaultGW
  -GW <default gw>
  ```

- To view the current network configuration, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network view-config
  ```

- To get the network time protocol (NTP), use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getNTP -I
  <Interface name>
  ```

- To set the NTP, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setNTP -I
  <Interface name> -N NTP1,NTP2,...
  ```

- To get the domain name system (DNS), use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getDNS -I
  <Interface name>
  ```

- To set the DNS, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setDNS -I
  <Interface name> -d DNS1,DNS2,...
  ```

- To get the IP address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getIP -I
  <Interface name>
  ```

- To set the IP address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network addIP -a
  <ADDRESS> \-gw <GATEWAY> -l <PREFIXLENGTH> -p <protocol type> -I <Interface
  name>
  ```

- To delete the IP address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network rmIP -I
  <Interface name> -a <ADDRESS>
  ```

- To enable a virtual local area network (VLAN), use the following command:

  ```
  openbmctool.py <connection options> network addVLAN -I <Interface name> -n
  <IDENTIFIER>
  ```

- To disable a virtual local area network (VLAN), use the following command:

  ```
  openbmctool.py <connection options> network deleteVLAN -I <Interface name>
  ```

- To view the DHCP configuration properties, use the following command:

  ```
  openbmctool.py <connection options> network viewDHCPConfig
  ```

- To configure the DHCP properties, use the following command:

  ```
  openbmctool.py <connection options> network configureDHCP -d <DNSENABLED> -n
  <HOSTNAMEENABLED> -t <NTPENABLED> -s <SENDHOSTNAMEENABLED>
  ```

  Note: DNSENABLED, HOSTNAMEENABLED, NTPENABLED, and SENDHOSTNAMEENABLED are boolean values (true or false).

- To reset the network settings to the factory defaults, use the following command:

  ```
  openbmctool.py <connection options> network nwReset
  ```

  Note: Reset settings are applied after the rebooting of the BMC.

Managing the system by using the IPMI
Learn how to configure and manage your system by using the Intelligent Platform Management Interface (IPMI).

Common IPMI commands
You can use **IPMI** commands to perform various managing tasks for your system.

Table 33. Common IPMI commands

Command option	**Description**
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `chassis power on`	Powers on the server.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `chassis power off`	Powers off the server.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `chassis status`	Checks the server status.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `chassis power cycle`	Power cycle the server.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `sol activate`	Activates SOL system console.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `sol deactivate`	Deactivates SOL system console.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `sel list`	Returns an error log.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `sdr list`	Lists status of all sensors.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `sol set retry-interval` *value*	Sets the default retry-interval value in milliseconds.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `fru print`	Prints the FRU information.
`ipmitool -I lanplus -H` *myserver.example.com* `-P` *mypass* `user list`	Lists the IPMI users.

Configuring the BMC IP address
Dynamic Host Configuration Protocol (DHCP) is the default network setup for the BMC in the 7063-CR2 HMC. To enable your network connection, you can connect to your system and use the Petitboot bootloader interface to configure the IP address of the BMC. If you do not plan to use DHCP, you can also set up a static IP address. Alternatively, you can use the HMC GUI to enable your network connection by navigating to **HMC Management** > **Console Settings** > **Change BMC/IPMI Network Settings**.

Before you begin

You must connect the network cable and a VGA monitor before you access the Petitboot bootloader interface.

If you encounter any problems in accessing the Petitboot bootloader interface, see Resolving a BMC access problem.

About this task

To use the Petitboot bootloader interface to set up or enable the network interface of the BMC, complete the following steps:

Procedure

1. Power on the server by pressing the power button on the front of the system. The system powers on to the Petitboot bootloader menu.

   Note: The boot process takes about 1 to 2 minutes to complete.

   When Petitboot loads, the monitor activates. Press any key to interrupt the boot process.
2. At the Petitboot bootloader main menu, select **Exit to Shell**.
3. Run the following command: `ipmitool lan print 2`. If this command returns an IP address, verify that is correct. To set a static IP address, follow these steps:

   Notes:

   - The following two LAN interfaces are available to BMC:
     - Shared interface is LAN1
     - Dedicated interface is LAN2
   - The `ipmitool lan print 2` command cannot display more than one IP address. To avoid this situation, you can set the static IP address to a different subnet from the default zero configuration networking IP address.

   a. Set the mode to static by running the following command: `ipmitool lan set 2 ipsrc static`.

   b. Set your IP address by running the following command: `ipmitool lan set 2 ipaddr ip_address`, where *ip_address* is the static IP address that you want to assign to this system.

   c. Set your netmask by running the following command: `ipmitool lan set 2 netmask netmask_address`, where *netmask_address* is the netmask for the system.

   d. Set your gateway server by running the following command: `ipmitool lan set 2 defgw ipaddr gateway_server`, where *gateway_server* is the gateway for this system.

   e. Confirm the IP address by running the following command: `ipmitool lan print 2`.

Performing a factory reset
Learn how to perform a factory reset on the system.

The factory reset function can take up to 15 minutes to complete. When the LED on the power button starts flashing, the system is ready to start again. Perform the factory reset with the host powered off. If you perform the factory reset while the host is running, the system shuts down immediately and restarts the BMC. If the BMC is on a static network, you must manually power on the system with the physical power button.

To perform a factory reset, run the following command:

```
ipmitool -I lanplus -U <username> -P <password> -H <BMC_IP or Hostname> raw 0x3A 0x11
```

Note: The system does not send a validation response. The following system output is normal:

```
Unable to send RAW command (channel=0x0 netfn=0x3a lun=0x0 cmd=0x11)
```

If you forgot the password of the BMC, you can run the following command while the host is running to perform a factory reset and to restore the default password:

```
ipmitool raw 0x3A 0x11
```

You must set up and configure the BMC IP address after performing the factory reset. For more information, see "Configuring the BMC IP address" on page 84.

Risks of using IPMI on IBM Power Systems and OpenPower Systems
Various risks that are associated with the Intelligent Platform Management Interface (IPMI) have been identified and documented in the information technology (IT) security community.

IBM Power Systems and OpenPower Systems provide IPMI access by default. A subset of these identified risks is applicable to IBM servers.

Note: Model 9080-M9S does not provide IPMI access by default.

Vulnerability Details

The IPMI service can become unresponsive after it receives and rejects multiple authentication attempts. You might receive a `insufficient resources for session` message if you use the IPMI immediately after the failed authentication attempts. This situation lasts for a few seconds and normal service is restored afterward.

Important: Repeated authentication failures can cause denial of service.

A list of common vulnerabilities and exposures (CVE) is listed in Table 34 on page 86.

Table 34. Common vulnerabilities and exposures

CVE ID	Description
CVE-2013-4037	The Remote Authenticated Key-Exchange Protocol (RAKP), which is specified by the IPMI standard for authentication, has flaws. Although the system does not allow the use of null passwords, a hacker might reverse engineer the RAKP transactions to determine a password. The authentication process for IPMI requires the management controller to send a hash of the requested password of the user to the client before the client authenticates. This process is a key part of the IPMI specification. The password hash can be broken by using an offline brute force or dictionary attack.
CVE-2013-4031	IBM Power Systems and OpenPower Systems are preconfigured with one IPMI user account, which has the same default login name and password on all affected systems. If a malicious user gains access to the IPMI interface by using this preconfigured account, the user can power off or on, or restart the host server, and create or change user accounts possibly preventing legitimate users from accessing the system. On OpenPower Systems, the default IPMI user name is `root`. Additionally, if a user fails to change the default user name and password on each of the systems that is deployed, the user has the same login information for each of those systems.
CVE-2013-4786	The IPMI 2.0 specification supports RMCP+ Authenticated Key-Exchange Protocol (RAKP) authentication, which allows remote attackers to obtain password hashes and conduct offline password guessing attacks by obtaining the hash-based message authentication code (HMAC) from a RAKP message 2 response from a BMC.

Configuration options and best practices

- Change the preconfigured user name and password when the server is deployed. This action prevents unauthorized users from gaining access to the system through the preconfigured user account.
- Do not disable the IPMI access for the user whose access credentials have been shared with the HMC via the **Console Inband Communication Credentials** task.

  Note: To launch the **Console Inband Communications Credentials** task, complete the following steps:

  1. In the navigation area, click the **HMC Management** icon, and then select **Console Settings**.
  2. In the content pane, click **Console Inband Communications Credentials**.

3. From the **Console Inband Communications Credentials** window, you can set the inland BMC credentials or modify an expired password for previously set inband BMC credentials for the HMC.

- If a user is not managing a server by using the IPMI, you can configure the system to disallow IPMI network access from the user accounts. This task can be accomplished by using the IPMItool utility or a similar utility for managing and configuring the IPMI management controllers. You can use the following IPMItool command to disable the network access for an IPMI user:

  ```
  ipmitool channel setaccess 1 #user_slot# privilege=15
  ```

  Note: Replace #user_slot# in the command with the actual slot number (1 - 12) and repeat for each configured user.

  This example shows the command when it is run directly on the server. If the IPMItool command is run remotely over the network, or if a different utility is used, the command might be different. See the documentation for the utility that you are using to determine the correct command syntax. Disallowing IPMI network access removes the ability to use the weakness that is present in the IPMI RAKP protocol to discover user account credentials.
- Use strong passwords that are at least 16 characters long with a mixture of upper and lowercase letters, numbers, and special characters. By using more complex passwords, it makes it more difficult for malicious users to discover valid user credentials.
- Keep the management network separate from the public network. Keeping the management network separate lessens security exposures by reducing the number of individuals who can access the systems.

Managing the system by using the OpenBMC GUI
Learn how to manage and configure your system by using the OpenBMC GUI.

Logging on to the OpenBMC GUI
Learn how to log on to the OpenBMC GUI.

To log on to the OpenBMC GUI, complete the following steps:

1. Open a supported web browser. In the address bar, enter the IP address of the BMC that you want to connect to. For example, you can use the format `https://<BMC IP>` in the address bar of the web browser.
2. From the **OpenBMC logon** window, enter the **Host** address of the BMC and the **Username** and **Password** that is assigned to you.

   Note: The default user ID is `root` and the default password is `0penBmc`.

   If you are using firmware level OP940.01, or later, the root password is expired by default. You must change the default password before you can access the BMC. For more information about changing the expired default password, see "Setting the password" on page 87.

   If you forgot your password, you can perform a factory reset of the system to restore the default password. To reset the system, see "Performing a factory reset" on page 85.
3. Click **Log in**.

Setting the password
Learn how to change and set the password for your **root** account and to help secure the system.

Improved BMC password policy

The baseboard management controller (BMC) **root** password must be set on first use for newly manufactured systems or after performing a factory reset of the system. This policy change helps to enforce that the BMC is not left in a state with a well-known password.

In firmware level OP940.01, and later, the root password is expired and must be changed before you can access the functions of the BMC. However, if you are upgrading the firmware level from a previous OpenBMC firmware level or if you are performing an operational installation, you do not have to change the password.

The default user ID is `root` and the default password is `0penBmc`. You can use the web application, the Redfish REST APIs, the OpenBMC tool command to change the password. You can also use the **Console Inband Communications Credentials** task in the HMC GUI to change the expired password.

Note: To launch the **Console Inband Communications Credentials** task, complete the following steps:

1. In the navigation area, click the **HMC Management** icon, and then select **Console Settings**.
2. In the content pane, click **Console Inband Communications Credentials**.
3. From the **Console Inband Communications Credentials** window, you can set the inland BMC credentials or modify an expired password for previously set inband BMC credentials for the HMC.

After changing the password, you can access the BMC with your usual interface. To change the password, you must first access the account with the correct credentials, and then use the password change function. If you attempt to access the BMC with an expired password, you must change the password before accessing other functions.

- To change your expired password by using the web interface, enter `https://<BMC_IP>` into a web browser and then enter the access credentials of the BMC. The web interface prompts you to enter a new password.
- To change your expired password through a network interface, you can use Redfish APIs. For instructions, see "Managing the system by using DMTF Redfish APIs" on page 96.
- To change your expired password by using the OpenBMC tool, run the `openbmctool set_password` subcommand. For example,

```
openbmctool.py -H <BMC IP address or BMC host name> -U <username> -P <password> set_password
-p <new password>
Attempting login...
200
User root has been logged out
```

  Where `200` is the response status that indicates success.

Note: The system might take up to 5 minutes to update the new password on the BMC. If you have trouble accessing your account, wait for 5 minutes and try again.

Also, with firmware level OP940.01, the BMC factory reset function resets the BMC password back to its default value and causes the default password to expire. This function means that after you perform the factory reset, you must change the password before you can access the BMC (even if you upgraded from an older firmware level).

To increase account security of the system, the administrator must complete the following steps:

1. Set a strong password for the root account. Strong passwords have at least 15 characters and include non-alphabetic characters. Initially, the password must not exceed 20 characters. Passwords can be changed later to a length greater than 20 characters, but IPMI access will be removed. Avoid using the **root** account, as the **root** account has more access to the BMC than an **Administrator** account. The root account can present a security risk if it is used incorrectly or maliciously. Use the root account only when it is required.
2. Create a separate account for each entity to manage the system. For example, you can create an **Administrator** account for yourself and for xCat, and create an **Operator** account for your staff. You can use the web interface or Redfish APIs to create a new account. When you create a new account, carefully consider which privilege role to assign to the user. Always use the least privilege role that is required.

   - To create a new account by using the web interface, see "Local users" on page 94.
   - To create a new account by using the Redfish APIs, see "Managing the system by using DMTF Redfish APIs" on page 96.

   If your BMC is using Lightweight Directory Access Protocol (LDAP), you can add users to the LDAP server.

3. Log off from the root account and switch to your personal **Administrator** account.

To increase the security of the system, the administrator can optionally configure access to the LDAP server. For more information, see “Basic commands and functionality of the OpenBMC tool ” on page 76.

Dashboard
The dashboard displays the overall information about the server and the BMC.

The following options are available on the title bar (located in the top portion of the dashboard):

- **Server information**: Displays the server name and BMC IP address.
- **Server health**: Displays the status of the server.
- **Server power**: Displays whether the server is powered on, powered off, or in an error state.
- **Date last refreshed**: Displays the date and time that the information was last refreshed. The time zone of the user is determined by the web browser.
- **Refresh**: Click **Refresh** to refresh the information.

The following menus are available from the menu pod (located in the left portion of the dashboard):

- **Server overview**
- **Server health**
- **Server control**
- **Server configuration**
- **Users**

Server overview
Learn about the options that are available from the **Server overview** task.

From the **Server overview** window, you can choose from any of the following available options:

- **Server information**: Displays the model, manufacturer, firmware version, and serial number of the server.
- **BMC information**: Displays the host name, BMC IP address, firmware version, and MAC address of the BMC.
- **Power information**: Displays the power consumption and power cap.
- **High priority events**: View any high priority events. Click **Refresh** to reload the information that is displayed here.
- **BMC time**: Displays the BMC time in the time zone of the user, which is determined by the web browser.
- **Turn on server LED**: Turn on or turn off the server LED.
- **Launch serial over LAN console**: Launches the Serial over LAN (SoL) console.
- **Edit network settings**: Edit the network settings.

Server health
Learn about the tasks that are available from the **Server health** menu.

From this menu, you can choose from any of the following available tasks:

Event log
View all events from the BMC.

Note: For descriptions and service actions for FQPSPxxxxxxx event codes, see Managing BMC-based systems by using the HMC (http://www.ibm.com/support/knowledgecenter/POWER9/p9ia7/p9ia7_kickoff.htm).

You can view and filter event log files from the BMC. From the **Event log** window, you can perform the following actions:

- Search through event logs by entering keywords and clicking **Search**.

- Filter the event logs by severity (**All**, **High**, **Medium**, or **Low**). You can select multiple severity levels.
- Filter the event logs by date range.
- Filter the event logs by event status (**All events**, **Resolved events**, and **Unresolved events**).
- Click any of the events that are listed to expand the event log file for more information. You can click **Copy** to copy the information to the clipboard.
- Select multiple event logs by clicking the checkbox next to event log. After you select the event logs, you can delete the logs by clicking **Delete** and then clicking **Yes** in the confirmation message. You can also mark event logs as read by clicking **Mark as resolved**.

Hardware status
View the hardware status and associated events of all hardware in the server.

You can view the hardware status of various hardware components in your server. Click any of the hardware components to expand the view for more information. You can search for specific hardware components by using the **Filter Hardware Components** search feature and then clicking **Filter**. You can also export the data by clicking **Export**.

Sensors
View all sensors that are present in the system.

You can view and filter sensors from the BMC. From the **Sensors** window, you can perform the following actions:

- Search and filter for specific sensors by using the **Search** feature and then clicking **Filter**.
- Filter sensors by severity (**All**, **Critical**, **Warning**, or **Normal**).
- Export the sensor data by clicking **Export**.

Server control
Learn about the tasks that are available from the **Server control** menu.

From this menu, you can choose from any of the following available tasks:

Server power operations
Learn how to view current server status and select power operations.

To update the **Host OS boot settings**, complete the following steps:

Note: It is not recommended to modify the **Host OS boot settings** for 7063-CR2 HMC, unless instructed by IBM Support.

1. Select the boot setting override type from the **Boot setting override** menu.
2. You can optionally select **Enable one time boot**.
3. Select whether to enable or disable the TPM policy by selecting **On** or **Off**.
4. Click **Save**.

To restart the server, complete the following steps:

1. Select the type of reboot from **Operations** > **Reboot server**:
   - Orderly: Operating system shuts down first and then the server reboots.
   - Immediate: Server reboots without the operating system shutting down. This might cause data corruption.
2. Click **Reboot**.

To shutdown the server, complete the following steps:

1. Select the type of shutdown from **Operations** > **Shutdown server**:
   - Orderly: Operating system shuts down first and then the server reboots.
   - Immediate: Server reboots without the operating system shutting down. This might cause data corruption.

2. Click **Shutdown**.

Manage power usage
Learn how to view the power consumption of the server and set a power cap.

Note: It is recommended to set the **power cap** to **Off** for 7063-CR2 HMC. By default, the **power cap** is set to **Off**.

To set a power cap, complete the following steps:

1. From the **Server power cap setting** section, set the **power cap** to **On**.
2. Specify the number of watts to keep the server power consumption at or below the specified value.
3. Click **Save settings**.

You can turn off the power cap by setting the **power cap** to **Off** and clicking **Save settings**.

Server LED
Learn how to turn on and turn off the server light-emitting diode (LED).

You can turn on or turn off the server LED by clicking the toggle switch to either **On** or **Off**.

Note: If the server has a liquid crystal display (LCD), you can use this control to display text (**On**) or not to display text (**Off**) on the LCD.

Reboot BMC
Learn how to restart the BMC and view the current BMC boot status.

Click **Reboot BMC** to restart the BMC.

Note: When you restart the BMC, your web browser looses connection with the BMC for several minutes. When the BMC is back online, you must log in again. If the **Log in** button is not available after you restart the BMC, close your web browser. Then, reopen the web browser and enter your BMC IP address.

Serial over LAN console
Learn how to view information over the serial port of the server.

You can launch the Serial over LAN (SoL) console that displays the output of the serial port of the server.

KVM
Learn how to launch the remote keyboard, video, and mouse (KVM) console.

You can launch the KVM console from this task and interact with the remote system.

Virtual media
Learn how to start a session by using a virtual media device.

To start a session, complete the following steps:

1. Under **Virtual media device**, click **Choose file**.
2. Select the file and click **Open**.
3. Click **Start** to start the session.

Server configuration
Learn about the tasks that are available from the **Server configuration** menu.

From this menu, you can choose from any of the following available tasks:

Network settings
Learn how to view and set common network, IPv4, and DNS settings.

To view network settings, select the **Network Interface** that you want to view. The **Hostname**, **MAC Address**, and **Default Gateway** are displayed under **Common settings**. The **DHCP setting**, **IPv4 IP addresses**, **Gateways**, and **Netmasks** are displayed under **IPv4 Settings**. Under **DNS settings**, all DNS servers are displayed.

To set network settings, complete the following steps:

1. Select the **Network Interface** that you want to set.
2. Edit the **Hostname**, **MAC Address**, or **Default Gateway** fields under **Common settings**.
3. Edit the **DHCP setting**, **IPv4 IP address**, **Gateway**, and **Netmask Prefix Length** under **IPv4 Settings**.
4. Click **Save Settings**.

Note: You can edit network settings only on firmware level OP920.01 or later.

To add an IPv4 address, complete the following steps:

1. Click **Add IPV4 address**.
2. Complete the **IPv4 IP address**, **Gateway**, and **Netmask Prefix Length** fields.
3. Click **Save Settings**.

To add a DNS address, complete the following steps:

1. Click **Add DNS Server**.
2. Enter the Internet Protocol (IP) address of the **DNS Server**.
3. Click **Save Settings**.

SNMP settings
Learn how to view and set the simple network management protocol (SNMP) with a hostname or Internet Protocol (IP) address and a port.

To set the SNMP, complete the following steps:

Note: Only SNMPv2 is supported.

1. Click **Add Manager**.
2. Enter the hostname or IP address and the port number.
3. Click **Save Settings**.

You can remove a manager by clicking the trash bin icon next to the manager that you want to remove.

Firmware
Learn how to manage the BMC and server firmware.

You can use the **BMC images** and **Server images** tables to manage firmware image files. The image file that is listed at the top, the image with the highest boot priority, is used the next time that the device is booted. You can change the boot order for the image file by clicking the arrow icons.

Learn about the different image states that are available:

- **Functional**: The running image on the device.
- **Active**: The image is available to boot from, but is not currently the running image. If the image is the top image in the relevant table, it becomes the functional image the next time the device is rebooted.
- **Activating**: The image is in the process of being activated and becomes either **Active** or **Failed**.
- **Failed**: The image failed to activate.
- **Ready**: The image is ready to be activated.
- **Invalid**: This image is an invalid image and cannot be activated.

The system firmware is a combination of the BMC firmware and the PNOR firmware. You must update both the BMC firmware and the PNOR firmware for the system to operate properly. If you update only one type of firmware, but do not update the other type of firmware, system errors might occur.

You can upload an image file from the workstation or from the Trivial File Transfer Protocol (TFTP) server. If you choose **Upload image file from workstation**, click **Choose a file** and specify the location of the image on the workstation storage device. Click **Upload** to upload the image file to the BMC server. If you choose **Download image file from TFTP server**, enter the TFTP server IP address in the **TFTP Server IP**

Address field and the file name in the **File Name** field. Click **Download** to download the image file to the BMC server.

After you load the new image file to the BMC server, you can activate the image file to make it available for use. Locate the image in the correct image table, and then click **Activate** > **Continue**. For a BMC image, an option of **Activate Firmware File Without Rebooting BMC** or **Activate Firmware File and Automatically Reboot BMC** is available. If you select **Activate Firmware File Without Rebooting BMC**, the BMC must be rebooted by using the **Reboot BMC** option to make the image become the **Functional** image. If you select **Activate Firmware File and Automatically Reboot BMC**, the BMC automatically reboots after the image is activated and the new image becomes the **Functional** image.

For a server image, after the image is activated, the server must be rebooted (or powered on if the server is off) for the image to become active. The **Reboot** (or **Power on**) option can be accessed from the **Server power operations** menu.

Date and time settings
Learn how to set the date and time.

To automatically set the date and time, complete the following steps:

1. Select **Obtain automatically from a network time protocol (NTP) server**.
2. Click **Add new NTP server**.
3. Enter the NTP server address.
4. Click **Save setttings**.

To manually add the date and time, complete the following steps:

1. Select **Manually set date and time**.
2. Enter the date and time.
3. Change the **Time owner** to the following values:
   - **BMC**: the BMC owns the time and can set the time.
   - **Host**: the host owns the time and can set the time.
   - **Split**: the BMC and the host own separate time.

     Note: Split is the recommended value for **Time owner** for 7063-CR2.
   - **Both**: both the BMC and the host can set the time.

Access control
Learn about the tasks that are available from the **Access control** menu.

From this menu, you can choose from any of the following available tasks:

LDAP
Learn how to configure lightweight directory access protocol (LDAP) settings and manage role groups.

LDAP authentication

To enable LDAP authentication, complete the following steps:

1. Select the **Enable LDAP authentication** checkbox.

   Note: If you want to secure LDAP by using Secure Sockets Layer (SSL), select the **Secure LDAP using SSL** checkbox. You must have a certificate authority (CA) and LDAP certificate for this function.
2. Select the service type as **Open LDAP** or **Active directory**.
3. Complete the required fields.
4. Click **Save**.

Role groups

To add a new role group, complete the following steps:

1. Click **Add role group**.
2. Enter a name for the role group.
3. Set the privilege of the role group to **Administrator**, **Operator**, **User**, or **Callback**.
4. Click **Save**.

To remove a new role group, complete the following steps:

1. Select the checkbox next to the role group or groups that you want to remove from the table.
2. Click **Remove role groups**.
3. Click **Remove** in the pop-up window.

To modify the privilege of a role group, complete the following steps:

Note: LDAP authentication must be enabled to modify group roles.

1. Select the checkbox next to the role group or groups that you want to modify from the table.
2. Click the **Edit** icon.
3. Change the privilege of the role group to **Administrator**, **Operator**, **User**, or **Callback**.
4. Click **Save**.

Local users
Learn how to add or remove new users, modify user settings, manage user account policy settings, and view privilege role descriptions.

To add a new user, complete the following steps:

1. Click **Add user**.
2. Set the account status to either **Enabled** or **Disabled**.
3. Enter a new username.

   Note: The username cannot start with a number. No special characters are allowed except for an underscore.
4. Set the privilege of the user to **Administrator**, **Operator**, **ReadOnly**, or **NoAccess (Callback)**.
5. Enter the password of the user.

   Note: Initially, the password must be in the range of 8 - 20 characters in length. Passwords can be changed later to a length greater than 20 characters, but IPMI access is removed. The system might take up to 5 minutes for the new password to update on the BMC. If you have trouble accessing your account, wait for 5 minutes and try again.
6. Reenter the password for confirmation.
7. Click **Add user**.

To remove a user, complete the following steps:

1. Click the checkbox next to the user or users that you want to remove from the table.
2. Click **Remove**.
3. Click **Remove** again in the pop-up window.

You can modify user settings by selecting the user from the table and clicking the edit icon. From the **Modify user** window, you can update the following properties:

- Account status: set to **Enabled** or **Disabled**.
- Username: change the name of the user.
- Privilege: change the account privileges of the user to **Administrator**, **Operator**, **ReadOnly**, or **NoAccess (Callback)**.

- User password: update the password of the user.

You can modify user account policy settings by clicking **Account policy settings**. From the **Account policy settings** window, you can update the following properties:

- Maximum failed login attempts: change the number of allowed failed login attempts.
- User unlock method: set to **Automatic after timeout** or **manual**.
- Only non-root accounts can be locked out.

You can view privilege role descriptions by clicking **View privilege role descriptions**.

Table 35. Privilege role descriptions

Role	Privileges	Guidance
Administrator	Can configure the BMC and manage users and sessions. Operational control over BMC functions.	Use this role for the most trusted users.
Operator	Operational control over BMC functions.	Use this role for users who manage routine operations.
ReadOnly	Read-only access to BMC functions.	Use this role for users who need to monitor the BMC, but do not need to operate it.
NoAccess (Callback)	No access to BMC functions.	Use this role for users who do not need access to the web interface or REST APIs.

SSL certificates
Learn how to generate a certificate signing request (CSR), add new certificates, and replace existing certificates.

Generating a CSR

To generate a new CSR, complete the following steps:

1. Click **Generate CSR**.
2. Complete the required fields under the **General** section.
3. Under **Private key** > **Key Pair Algorithm**, select the algorithm as **EC** or **RSA**.
4. Click **Generate CSR**.

Certificates

To add a new certificate, complete the following steps:

1. Click **Add new certificate**.
2. Select the certificate type as **HTTPS Certificate**, **LDAP Certificate**, or **CA Certificate**.
3. Click **Choose file** to select the certificate.
4. Click **Open**.
5. Click **Save**.

To replace certificates, complete the following steps:

1. Select the certificate that you want to replace from the table.
2. Click the refresh icon.
3. Click **Choose file** to select the new certificate.

4. Click **Open**.
5. Click **Replace**.

Managing the system by using DMTF Redfish APIs
OpenBMC-based systems can be managed by using the DMTF Redfish APIs.

Overview

Redfish is a REST API used for platform management and is standardized by the Distributed Management Task Force, Inc. (http://www.dmtf.org/standards/redfish).

Redfish enables platform management tasks to be controlled by client scripts that are developed by using secure and modern programming paradigms.

The Redfish API enables provisioning of tunable parameters for better utilization of power.

IBM OpenBMC-based systems support DMTF Redfish API (DSP0266, version 1.7.0, published on 20 May 2019) for systems management.

A copy of the Redfish schema files that are in JSON format are published by DMTF (http://redfish.dmtf.org/schemas/v1/) and are packaged in the firmware image.

The schema files that are distributed in the chip enable proper functioning of the APIs in deployments that have no wide area network (WAN) connectivity.

Note: The Redfish API is enabled by default and the Redfish service cannot be enabled or disabled by the user.

Firmware levels

Redfish APIs are supported on OpenPOWER (OP) firmware level OP940, or later.

Communication prerequisites for Redfish on OpenBMC-based servers

Depending on the current firmware level and network deployment, complete the following prerequisite tasks:

- Upgrade the server firmware level to OP940, or later.
- Identify the IP address of the BMC.
- Install and run cURL (https://curl.haxx.se/) with the method, Uniform Resource Indicator (URI), and the request body as parameters to communicate with the Redfish service.
- Install Python on the client system (typically a Linux host).
- Optionally, install and run DMTF Redfishtool (https://github.com/DMTF/Redfishtool).

Interacting with the Redfish service

To interact with the Redfish service, complete the following steps.

1. Create an authenticated login session (POST method on the `/redfish/v1/SessionService/Sessions` resource).
2. Extract and save the following details:
   - Authentication token (found in the **X-Auth-Token** header of the response)
   - Session URI (found in the **Location** header of the response)
3. To read the properties of a resource, send a **GET** request with the **X-Auth-Token** header for the URI of the resource.
4. To set a property of a resource, send a **PATCH** request with the **X-Auth-Token** header for the URI of the resource, the property name, type, and value encoded as a JSON body.
5. Extract and parse the response from the Redfish service that contains the JSON body.

Redfish service home page URI

The Redfish service home page URI (also known as the service ROOT) can be accessed by retrieving the URI: `https://<ip:port>/redfish/v1`. The response to this URI is a high-level site map that enables a traversal of the Redfish service by using a hypermedia API paradigm.

Interpreting the data returned by the Redfish service

The format and structure of the data is defined in the schema files. Schema files are JSON files that describe the data that is sent by the Redfish service. You can use the schema files to understand the data that is sent by the Redfish service and to validate the response that is sent by the Redfish service.

Location of the schema files

DMTF publishes the schema files for the standard data that is used in Redfish.

The Redfish schema files in JSON format are hosted in the DMTF schema repository at http://redfish.dmtf.org/schemas/v1/

Supported schema files

The following schema files are supported for OpenBMC-based systems:

- Account
- AccountCollection
- AccountService
- Certificate
- CertificateCollection
- CertificateLocations
- CertificateService
- Chassis
- ChassisCollection
- ComputerSystem
- ComputerSystemCollection
- EthernetInterface
- EthernetInterfaceCollection
- LogEntry
- LogService
- LogServiceCollection
- Manager
- ManagerCollection
- ManagerNetworkProtocol
- Memory
- MemoryCollection
- Processor
- ProcessorCollection
- Role
- RoleCollection
- ServiceRoot
- Session

- SessionCollection
- SessionService
- SoftwareInventory
- SoftwareInventoryCollection
- ThermalPower
- UpdateService

Accessing the common system management functions on the Redfish service by using cURL command

The following examples show the client URL (cURL) commands that can be used to access the common functions that are supported by the OpenBMC Redfish APIs:

Note: In all cURL commands, *${BMC}* is the IP address of the BMC.

- To view major collections, run the following commands:
  - Chassis collection:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Chassis
    ```

  - Manager collection:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Managers
    ```

  - System collection:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Systems
    ```

- To view the chassis, manager, and system resources, run the following commands:
  - Chassis resource:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Chassis/chassis
    ```

  - Manager resource:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Managers/bmc
    ```

  - System resource:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Systems/system
    ```

- To perform host power control operations, run the following commands:
  - Host power on:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "On"}'
    ```

  - Host soft power off:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "GracefulShutdown"}'
    ```

  - Host hard power off:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "ForceOff"}'
    ```

  - Restart host:

```
-X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
'{"ResetType": "GracefulRestart"}'
```

- To view the host power control resource, run the following command:

  ```
  curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/Actions/
  ```

- To view the log resource, run the following command:

  ```
  curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/LogServices/EventLog/
  Entries
  ```

- To view sensor resources, run the following commands:
  - Power resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/chassis/Power
    ```

  - Thermal resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/chassis/Thermal
    ```

  - Sensor resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/chassis/Sensors
    ```

- To view inventory resources, run the following commands:
  - Memory resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/Memory
    ```

  - Processor resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/Processors
    ```

  - Power supply 0 resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/powersupply0
    ```

  - Power supply 1 resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/powersupply1
    ```

  - Motherboard resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/motherboard
    ```

- To update the firmware, run the following commands:
  - By using an image file from your system:

    ```
    curl -u root:0penBmc -curl k -s  -H "Content-Type: application/octet-stream" -X POST -T
    <image file path> https://${BMC}/redfish/v1/UpdateService
    ```

  - By using a Trivial File Transfer Protocol (TFTP) server:

    ```
    curl -u root:0penBmc -k -s -d '{"ImageURI":"<TFTP IP Address>/<File name on
    TFTP server>","TransferProtocol":"TFTP"}' -X POST https://${BMC}/redfish/v1/UpdateService/
    Actions/UpdateService.SimpleUpdate
    ```

- To create a new local account, run the following command:

  -
    ```
    curl -X POST https://${BMC}/redfish/v1/AccountService/Accounts/ -d '{"UserName": "admin",
    "Password": "NEWPASSWORD", "RoleId": "Administrator"}'
    ```

Where `admin` is the name of the user that you want to create, `NEWPASSWORD` is the new password, and `RoleId` maps to the privilege role.

- To change the account password, run the following command:

  -
    ```
    curl -X POST https://${BMC}/redfish/v1/AccountService/Accounts/root -d '{"Password":
    "NEWPASSWORD"}'
    ```

  Where `root` is the account name or user ID and `NEWPASSWORD` is the new password.

For more information about selecting a username, password, or role, see "Local users" on page 94.

Configure BMC connectivity (7063-CR2)
You can configure or view the network settings on the BMC for the management console.

Notes:

- This task applies only to the 7063-CR2. This connection is required to access the baseboard management controller (BMC) on the HMC.
- The settings in this task are applicable only to the dedicated IPMI or BMC network port.

To configure the BMC connection, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change BMC/IPMI network settings**.
3. Select the connection mode (**DHCP** or **Static**).

   If you select **Static** mode, complete the following addresses:

   - **IP address**
   - **Subnet mask**
   - **Gateway**
4. Click **OK**.

You can also configure the BMC network connection by using the Petitboot bootloader interface. For more information, see Configuring the firmware IP address.

Configure BMC connectivity (7063-CR1)
You can configure or view the network settings on the BMC for the management console.

Note: This task applies only to the 7063-CR1. This connection is required to access the baseboard management controller (BMC) on the HMC.

To configure the BMC connection, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change BMC/IPMI network settings**.
3. Select the connection mode (**DHCP** or **Static**).

   If you select **Static** mode, complete the following addresses:

   - **IP address**
   - **Subnet mask**
   - **Gateway**
4. Click **OK**.

You can also configure the BMC network connection by using the Petitboot bootloader interface. For more information, see Configuring the firmware IP address.

Setting the IPv4 address
Learn how to set your IPv4 address on the HMC.

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **Basic Settings** tab.
6. Select an IPv4 address.
7. If you selected to specify an IP address, enter the TCP/IP interface address and the TCP/IP interface network mask.
8. Click **OK**.

Setting the IPv6 address
Learn how to set your IPv6 address on the HMC.

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **IPv6 Settings** tab.
6. Select an **Autoconfig** option or add a static IP address.
7. If you added an IP address, enter the IPv6 address and the prefix length and click **OK**.
8. Click **OK**.

Using only IPv6 addresses
Learn how to configure the HMC so that it uses only IPv6 addresses.

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Select **No IPv4 address**.
6. Click the **IPv6 Settings** tab.
7. Select **Use DHCPv6 to configure IP settings** or add static IP addresses, then click **OK**.

What to do next

After you click **OK**, you must restart your HMC for these changes to take effect.

Changing HMC firewall settings

In an open network, a firewall is used to control outside access to your company network. The HMC also has a firewall on each of its Ethernet adapters. To control the HMC remotely or give remote access to others, modify the firewall settings of the Ethernet adapter on the HMC that is connected to your open network.

About this task

To configure a firewall, use the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **Firewall** tab.
6. Using one of the following methods, you can allow any IP address by using a particular application through the firewall, or you can specify one or more IP addresses:
   - Allow any IP address by using a particular application through the firewall:
     a. From the top box, highlight the application.
     b. Click **Allow Incoming**. The application displays in the bottom box to signify that it is selected.
   - Specify which IP addresses to allow through the firewall:
     a. From the top box, highlight an application.
     b. Click **Allow Incoming by IP Address**.
     c. On the Hosts Allowed window, enter the IP address and the network mask.
     d. Click **Add** and click **OK**.
7. Click **OK**.

   Notes:

   - For more information about enabling remote restricted shell access, see "Enabling remote restricted shell access" on page 102.
   - For more information about enabling remote web access, see "Enabling remote web access" on page 103.

Enabling remote restricted shell access

You can enable remote restricted shell access when you configure a firewall.

About this task

To enable remote restricted shell access, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security** > **Systems and Console Security**.
2. Click **Enable Remote Command Execution**.
3. Select **Enable remote command execution using the ssh facility** and then click **OK**.

What to do next

Now remote restricted shell access is enabled.

Enabling remote web access

You can enable remote web access to your Hardware Management Console (HMC).

About this task

To enable remote web access, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security** > **Systems and Console Security**.
2. Click **Enable Remote Command Execution**.
3. Select **Enable** and then click **OK**.

What to do next

Now remote web access is enabled.

Configuring a routing entry as the default gateway

Learn how to configure a routing entry as the default gateway. This task is available when you are using an open network.

Before you begin

To configure a routing entry as the default gateway, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Customize Network Settings window opens.
3. Click the **Routing** tab.
4. In the Default gateway information section, enter the gateway address and gateway device of the routing entry you want to set as the default gateway.
5. Click **OK**.

Configuring domain name services

If you plan to set up an open network, configure domain name services.

About this task

If you plan to set up an open network, configure domain name services. Domain Name System (DNS) is a distributed database system for managing host names and their associated Internet Protocol (IP) addresses. Configuring domain name services includes enabling DNS and specifying the domain suffix search order.

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Change Network Settings window opens.
3. Click the **Name Services** tab.
4. Select **DNS enabled** to enable DNS.
5. Specify the DNS server and domain suffix search order and click **Add**.
6. Click **OK**.

Configuring domain suffixes

The list of domain suffixes is used to resolve an IP address that starts with the first entry in the list.

About this task

The domain suffix is a string that is appended to a host name that is used to help resolve its IP address. For example, a host name of `myname` might not be resolved. However, if the string `myloc.mycompany.com` is an element in the domain suffix table, then an attempt is made to resolve `myname.mloc.mycompany.com`.

To configure a domain suffix entry, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Customize Network Settings window opens.
3. Click the **Name Services** tab.
4. Enter a string to be used as a domain suffix entry.
5. Click **Add** to add it to the list.

Configuring the HMC so that it uses LDAP remote authentication

You can configure your Hardware Management Console (HMC) so that it uses LDAP (Lightweight Directory Access Protocol) remote authentication.

Before you begin

When a user logs in to the HMC, authentication is first performed against a local password file. If a local password file is not found, the HMC can contact a remote LDAP server for authentication. You must configure your HMC so that it uses LDAP remote authentication.

Note: Before you configure the HMC so that it uses LDAP authentication, you must ensure that a working network connection exists between the HMC and the LDAP servers. For more information about configuring HMC network connections, see "Configuring the HMC network types" on page 70.

About this task

To configure your HMC so that it uses LDAP authentication, complete the following steps:

Procedure

1. In the navigation area, click the **Users and Security** icon , and then select **Systems and Console Security**.
2. In the content pane, select **Manage LDAP**. The LDAP Server Definition window opens.
3. Select **Enable LDAP**.
4. Define an LDAP server to use for authentication.
5. Define the LDAP attribute that is used to identify the user that is being authenticated. The default is **uid**, but you can use your own attributes.
6. Define the distinguished name tree, also known as the search base, for the LDAP server.
7. Click **OK**.
8. If a user wants to use LDAP authentication, the user must configure their profile so that it uses LDAP remote authentication instead of local authentication.

Configuring the HMC so that it uses Key Distribution Center servers for Kerberos remote authentication

You can configure the HMC so that it uses Key Distribution Center (KDC) servers for Kerberos remote authentication.

Before you begin

When a user logs in to the HMC, authentication is first verifies against a local password file. If a local password file is not found, the HMC can contact a remote Kerberos server for authentication. You must configure your HMC so that it uses Kerberos remote authentication.

Note: Before you configure the HMC so that it uses KDC servers for Kerberos remote authentication, you must ensure that a working network connection exists between the HMC and the KDC servers. For more information about configuring HMC network connections, see "Configuring the HMC network types" on page 70.

About this task

To configure the HMC so that it uses KDC servers for Kerberos remote authentication, complete the following steps:

Procedure

1. Enable the Network Time Protocol (NTP) service on the HMC and set the HMC and the KDC servers to synchronize time with the same NTP server. To enable the NTP service on the HMC, complete the following steps:
   a) In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
   b) In the content pane, select **Change Date and Time**.
   c) Select the **NTP Configuration** tab.
   d) Select **Enable NTP service on this HMC**.
   e) Click **OK**.
2. Configure each remote HMC user's profile so that it uses Kerberos remote authentication instead of local authentication.
3. Optionally, you can import a service-key file into this HMC. The service-key file contains the host principal that identifies the HMC to the KDC server. Service-key files are also known as *keytabs*. To import a service-key file into this HMC, complete the following steps:
   a) In the navigation area, click the **Users and Security** icon , and then select **Systems and Console Security**.
   b) In the content pane, select **Manage KDC**.
   c) Select **Actions > Import Service Key**. The Import Service Key window opens.
   d) Type the location of the service key file.
   e) Click **OK**.
4. Add a new KDC server to this HMC. To add a new KDC server to this HMC, complete the following steps:
   a) In the navigation area, click the **Users and Security** icon , and then select **Systems and Console Security**.
   b) In the content pane, select **Manage KDC**.

c) Select **Actions > Add KDC Server**. The Import Service Key window opens.

d) Type the realm and the host name or IP address of the KDC server.

e) Click **OK**.

Configuring the local console to report errors to service and support

Configure this HMC so that it can call-home errors by using LAN connectivity.

Configuring the HMC so that it can connect to service and support by using the call-home setup wizard
Configure the HMC so that it is a call-home server by using the call-home wizard.

Before you begin

This procedure describes how to configure the HMC as a call-home server by using direct (LAN-based) and indirect (SSL) connections to the internet.

Before you begin this task, ensure that:

- The network administrator verifies that connectivity is allowed. For more information, see "Preparing for HMC configuration" on page 59.
- If you are configuring internet support through a proxy server, you must also have the following information:
  - The IP address and port of the proxy server
  - The proxy authentication information
- The adapter that is designated as **eth1** (the one that is designated as an open network) is used. For more information, see "Choosing network settings on the HMC" on page 51.
- An Ethernet cable physically connects the HMC to the LAN.

To configure the HMC so that it is a call-home server by using the call-home wizard, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon, and then select **Service Management**.
2. In the content pane, click **Call-Home Setup Wizard**. The Connectivity and Call-Home Servers wizard opens. Follow the instructions in the wizard to configure call-home.

Configuring the local console to report errors to service and support
Configure this HMC so that it can call-home errors by using LAN connectivity.

Configuring an HMC to contact service and support by using LAN-based internet and SSL
Describes how to configure the HMC as a call-home server by using direct (LAN-based) and indirect (SSL) connections to the internet.

Before you begin

Before you begin this task, ensure that:

- The network administrator verifies that connectivity is allowed. For more information, see "Preparing for HMC configuration" on page 59.
- Customer contact information is configured. Verify the contact information by going to the HMC interface and clicking **Serviceability>Service Management > Manage Customer Information**.
- If you are configuring internet support through a proxy server, you must also have the following information:
  - The IP address and port of the proxy server

  - The proxy authentication information
- You need at least one open network interface configured. For more information, see “Private and open networks in the HMC environment” on page 53.
- An Ethernet cable physically connects the HMC to the LAN.

About this task

To configure the HMC as a Call Home server by using LAN-based internet and SSL, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the Connectivity section, click **Manage Outbound Connectivity**. The Call-home Server Consoles window opens.
3. Click **Configure.**
4. In the Outbound Connectivity Settings window, check **Enable local system as call-home server**.
5. Accept the agreement.
6. In the Outbound Connectivity Settings window, select the **Internet** page.
7. Check the **Allow an existing internet connections for service** box.
8. If you are using an SSL proxy, check the **Use SSL proxy** box.
9. If you are using an SSL proxy, complete the proxy's address and port. Obtain this information from the network administrator.
10. If you checked **Use SSL proxy** and the proxy requires user ID and password authentication, check the **Authenticate with the SSL proxy** box. Type the user ID and password. Obtain the user ID and password from the network administrator.
11. Select the **Protocol to Internet** you want to use.
12. On the **Internet** page, click **Test**.
13. In the Test internet window, click **Start**.
14. Verify that the test completes successfully.
15. In the Test internet window, click **Cancel**.
16. In the Outbound Connectivity Settings window, click **OK**.

Choosing existing call-home servers to connect to service and support for this HMC
Choose existing Hardware Management Console (HMC) call-home servers that are recognized or discovered by the HMC to report errors.

Before you begin

Discovered HMCs are HMCs that are enabled as call-home servers and are either on the same subnet or manage the same managed system as this HMC.

To choose a discovered HMC to call home when the HMC reports errors, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Manage Outbound Connectivity**. The Call-Home Server Consoles window opens.

3. Click **Use discovered call-home server consoles**. The HMC displays the IP address or host name of the HMCs configured for call-home.
4. Click **OK**.

Results

You can also manually add existing HMC call-home servers that are on a different subnet. Select the IP address or host name of the HMC that is configured for call home and click **Add** and then click **OK**.

Verifying that your connection to service and support is working
Test problem reporting to ensure that connection to service and support is working.

About this task

To verify that your call-home configuration is working, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Create Event**.
3. Select **Test Automatic problem Reporting** and type a comment.
4. Click **Request Service**. Wait a few minutes for the request to be sent.
5. In the Service Management window, select **Manage Events**.
6. Select **All open problems**.
7. Verify that a PMH event and number is assigned to the problem number you opened.
8. Select that event and click **Close**.
9. On the **Close** window, type your name and a brief comment.

Authorizing users to view collected system data
You must authorize users to view data about your systems.

Before you begin

Before you authorize users to view collected system data, you must obtain an IBM ID. For more information about obtaining an IBM ID, see "Preinstallation configuration worksheet for the HMC" on page 60.

About this task

To authorize users to view collected system data, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, select **Authorize User**.
3. Enter your IBM ID.
4. Click **OK**.

Transmitting service information
You can transmit information to your service provider immediately, or you can schedule the information to be sent regularly.

Before you begin

IBM provides personalized web functions that use information that is collected by IBM Electronic Service Agent. To use these functions, you must first register on the IBM Registration website at http://www.ibm.com/account/profile. To authorize users to use the Electronic Service Agent information to personalize the web functions, see "Authorizing users to view collected system data" on page 108. For more information about the benefits of registering an IBM ID with your systems, see http://www.ibm.com/support/electronic.

Note: You must transmit service provider information as soon as the HMC is installed and configured for use.

About this task

To transmit service information, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon, and then select **Service Management**.
2. In the content pane, click **Transmit Service Information.**
3. Complete the tasks in the **Transmit Service Information** window, and click **OK**.

Configuring the Events Manager for Call Home

Learn how to configure the Events Manager for Call Home task. You can monitor and approve any data that is being transmitted from an HMC to IBM through this task.

The Events Manager for Call Home mode (enabled or disabled) is set by using the HMC command line interface. Enabling the Events Manager for Call Home task blocks the HMC from automatically calling home events as they occur. To prevent events that are called home without approval, all HMCs running in this environment must have the Events Manager for Call Home enabled.

To enable or disable the Events Manager for Call Home task, run the following command:

```
chhmc -c emch

-s {enable | disable}

[--callhome {enable | disable}]

[--help]
```

Note: Enabling the Events Manager for Call Home task holds call home events until they are approved for the call home task. If you disable the Events Manager for Call Home task, it does not automatically enable the call home feature. This setup prevents any unintended call home of data back to IBM. Choose from the following command options to set up the required configuration:

- To enable the Events Manager for Call Home task: **`chhmc -c emch -s enable`**
- To disable the Events Manager for Call Home task and to re-enable automatic call home: **`chhmc -c emch -s disable --callhome enable`**
- To disable the Events Manager for Call Home task and not re-enable automatic call home: **`chhmc -c emch -s disable --callhome disable`**

Ensure that the HMC can communicate with other HMCs deployed in this environment. The Events Manager for Call Home has a test connection function when an HMC is registered.

You can register the HMC with the Events Manager for Call Home. After you register the HMC, the events manager queries the registered HMC for any events that are waiting to be called home to IBM. The Events

Manager shows what data is being sent back to IBM and approves these events. After approval, the Event Manager notifies the registered HMC that it can proceed with the call home operation.

The Events Manager for Call Home task can be run from any HMC or from multiple HMCs. To register a management console with the Events Manager for Call Home task, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Events Manager for Call Home**.
2. From the **Events Manager for Call Home** pane, click **Manage Consoles**.
3. From the **Manage Registered Consoles** window, click **Add Console** to enter information to register a management console with the Events Manager for Call Home task.
4. Click **OK** to commit the changes to the list of registered management console.

Note: The Events Manager for Call Home can be used with the event manager mode disabled. You can still register the HMC and view events in the events manager, but Events Manager does not control when the events are called home.

Setting passwords for the managed system

You must set passwords for both your server and Advanced System Management (ASM). Read more about how to use the HMC interface to set these passwords.

Before you begin

If you received the message `Authentication Pending`, the HMC prompts you to set the passwords for the managed system.

About this task

If you did not receive the message Authentication Pending, complete the following steps to set the passwords for the managed system.

Updating your server password

Before you begin

To update your server password, complete the following steps:

Procedure

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. Click **Change Password**. The Update Password window opens.
3. Type the required information and click **OK**.

Updating your Advanced System Management (ASM) general password

Before you begin

Note: The default password for the general user ID is `general`, and the default password for the administrator ID is `admin`.

To update your ASM general password, complete the following steps:

Procedure

1. In the navigation area of the HMC, select the managed system.

2. In the Tasks area, click **Operations**.
3. Click **Advanced System Management (ASM)**. The Launch ASM Interface window opens.
4. Select a Service Processor IP Address and click **OK**. The ASM interface opens.
5. On the ASMI Welcome pane, specify your user ID and password, and click **Log In**.
6. In the navigation area, expand **Login Profile**.
7. Select **Change Password**.
8. Specify the required information, and click **Continue**.

Resetting the Advanced System Management (ASM) administrator password

Before you begin

To reset the administrator password, contact an authorized service provider.

Testing the connection between the HMC and the managed system

Learn how to verify that you are properly connected to the network.

About this task

To test the network connectivity, you must be a member of one of the following roles:

- Super administrator
- Service representative

To test the connection between the HMC and the managed system, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Test Network Connectivity**.
3. In the Ping tab, type the host name or IP address of any system to which you want to connect. To test an open network, type the gateway. Click **Ping**.

Results

If you have not created any logical partitions, you cannot ping the addresses. You can use the HMC to create logical partitions on your server. For more information, see Logical partitioning.

To understand how the HMC can be used in a network, see "HMC network connections" on page 51.

For more information about configuring the HMC to connect to a network, see "Configuring the HMC by using the menus " on page 67.

Postconfiguration steps

After you install and configure the HMC, back up HMC data as necessary.

Backing up management console data

This task backs up (or archives) the data that is stored on your HMC hard disk that is critical to support HMC operations.

Before you begin

Your remote system must have Network File System (NFS) or Secure Shell (ssh) configured, and this network must be accessible from the HMC. To complete this task, you must shut down and reboot the HMC. Use only the HMC to perform these tasks.

About this task

To back up the HMC hard disk drive to a remote system, you must be a member of one of the following roles:

- Super administrator
- Operator
- Service representative

Back up the HMC data after changes have been made to the HMC or information associated with logical partitions.

The HMC data stored on the HMC hard drive can be saved to a DVD-RAM on a local system, a remote system mounted to the HMC file system (such as NFS), or sent to a remote site using File Transfer Protocol (FTP).

Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

Using the HMC, you can back up all important data, such as the following:

- User-preference files
- User information
- HMC platform-configuration files
- HMC log files
- HMC updates through Install Corrective Service.

To back up the HMC hard drive to a remote system, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Backup Management Console Data**.
3. From the **Backup Management Console Data** window, select the archive option you want to perform.
4. Click **Next**, then follow the appropriate instructions depending on the option you chose.
5. Click **OK** to continue with the backup process.

Updating, upgrading, and migrating your HMC machine code

Updates and upgrades are periodically released for the HMC to add new functionality and to improve existing features. Learn more about the differences between updating, upgrading, and migrating your HMC machine code. Also learn how to perform an HMC machine code update, upgrade, or migration.

When you are finished with each of these tasks, the HMC reboots but the partitions do not.

Updating HMC code

Applies maintenance to an existing HMC level

Does not require that you perform the **Save upgrade data** task

Upgrading HMC code

Replaces HMC software with a new release or fix level of the same program

Requires that you boot from recovery media

Migrating HMC code

Moves HMC data from one HMC version to another

A migration is a type of upgrade.

Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

Determining your HMC machine code version and release

Find out how to view the HMC machine code version and release.

About this task

The level of machine code on the HMC determines the available features, including concurrent server firmware maintenance and enhancements to upgrading to a new release.

To view the HMC machine code version and release, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the **Current HMC Driver Information** heading, including: the HMC version, release, maintenance level, build level, and base versions.

Obtaining and applying machine code updates for the HMC with an Internet connection

Learn how to obtain machine code updates for the HMC when the HMC has an Internet connection.

About this task

To obtain machine code updates for the HMC, complete all steps.

Step 1. Ensure that you have an Internet connection

About this task

To download updates from the service and support system or website to your HMC or server, you must have one of the following connections:

- SSL connectivity with or without a SSL proxy
- Internet VPN

To ensure that you have an Internet connection, do the following:

Procedure

1. In the navigation area, click the **Serviceability** icon, and then select **Service Management**.
2. In the content pane, **Manage Outbound Connectivity**.
3. Select the tab for the type of outbound connectivity that you chose for your HMC (Internet VPN or SSL connectivity).

   Note: If a connection to service and support does not exist, set up the service connection before proceeding with this procedure. For instructions on how to set up a connection to service and support, see Setting up your server to connect to IBM service and support.
4. Click **Test**.
5. Verify that the test completes successfully.

   If the test is not successful, troubleshoot your connectivity and correct the problem before proceeding with this procedure. Alternatively, you can obtain the update on DVD.

   Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

6. Continue with “Step 2. View the existing HMC machine code level” on page 114.

Step 2. View the existing HMC machine code level

About this task

To view the existing HMC machine code level, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Continue with “Step 3. View the available HMC machine code levels” on page 114.

Step 3. View the available HMC machine code levels

About this task

To view the available HMC machine code levels, complete the following steps:

Procedure

1. From a computer or server with an Internet connection, go to http://www.ibm.com/eserver/support/fixes.
2. Select the appropriate family in the Product family list.
3. Select **Hardware Management Console** in the Product or fix type list.
4. Click **Continue**.

   The Hardware Management Console site is displayed.
5. Scroll down to your HMC Version level to view available HMC levels.

   Note: If you prefer, you can contact service and support.
6. Continue with “Step 4. Apply the HMC machine code update” on page 114.

Step 4. Apply the HMC machine code update

About this task

To apply the HMC machine code update, complete the following steps:

Procedure

1. Before you install updates to the HMC machine code, back up critical console information on your HMC.

   For instructions, see “Backing up management console data” on page 111. Then continue with the next step.

2. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.
3. In the content pane, click **Update the Hardware Management Console**. The Install Corrective Service Wizard opens.
4. Follow the instructions in the Wizard to install the update.
5. Shut down and then restart the HMC for the update to take effect.
6. Click **Log on and launch the Hardware Management Console web application**.
7. Log in to the HMC interface.

Step 5. Verify that the HMC machine code update installed successfully

About this task

To verify that the HMC machine code update installed correctly, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Verify that the version and release match the update that you installed.
5. If the level of code displayed is not the level that you installed, perform the following steps:

   a. Select the network connection on the HMC.

   b. Retry the firmware update using a different repository.

   c. If the problem persists, contact your next level of support.

Obtaining and applying machine code updates for the HMC using DVD or an FTP server

Learn how to obtain machine code updates for the Hardware Management Console (HMC) by using DVD or an FTP server.

About this task

To obtain HMC machine code updates, complete all steps.

Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

Step 1. View the existing HMC machine code level

Before you begin

To view the existing HMC machine code level, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Continue with "Step 2. View the available HMC machine code levels" on page 116.

Step 2. View the available HMC machine code levels

Before you begin

To view the available HMC machine code levels, complete the following steps:

About this task

Procedure

1. From a computer or server with an internet connection, go to the Fix Central website.
2. Scroll down to your HMC Version level to view available HMC levels.

   Note: If you prefer, you can contact IBM service and support.
3. Continue with "Step 3. Obtain the HMC machine code update" on page 116.

Step 3. Obtain the HMC machine code update

Before you begin

To obtain the HMC machine code update, complete the following steps:

About this task

You can order the HMC machine code update through the Fix Central website, contact service and support, or download it to an FTP server.

Ordering the HMC machine code update through the Fix Central website

1. From a computer or server with an Internet connection, go to the Fix Central website.
2. Under Supported HMC products, select the latest HMC level.
3. Scroll down to the File names / Package area and locate the update that you want to order.
4. In the Order column, select **Go**.
5. Click **Continue** to sign in with your IBM ID.
6. Follow the on-screen prompts to submit your order.

Downloading the HMC machine code update to removable media

1. From a computer or server with an Internet connection, go to Fix Central website.
2. Under Supported HMC products, select the latest HMC level.
3. Scroll down to the File names / Package area and locate the update that you want to download.
4. Click the update that you want to download.
5. Accept the license agreement, and save the update to your removable media.

What to do next

When you are finished, continue with “Step 4. Apply the HMC machine code update” on page 117.

Step 4. Apply the HMC machine code update

Before you begin

To apply the HMC machine code update, complete the following steps:

Procedure

1. Before you install updates to the HMC machine code, back up HMC data. For more information, see “Backing up management console data” on page 111.
2. If you obtained or created the update on DVD-RAM, insert it into the DVD drive on the HMC. If you obtained or created the update on a USB memory device, insert the memory device.
3. Before you install updates to the HMC machine code, back up critical console information on your HMC.

   For instructions, see “Backing up management console data” on page 111. Then continue with the next step.
4. In the navigation area, click the **HMC Management** icon, and then select **Console Management**.
5. In the content pane, click **Update the Hardware Management Console**. The Install Corrective Service Wizard opens.
6. Follow the instructions in the Wizard to install the update.
7. Shut down, restart, and log back in to the HMC for the update to take effect.
8. Continue with “Step 5. Verify that the HMC machine code update installed successfully” on page 117.

Step 5. Verify that the HMC machine code update installed successfully

Before you begin

To verify that the HMC machine code update installed successfully, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Verify that the version and release match the update that you installed.
5. If the level of code that is displayed is not the level that you installed, perform the following steps:

   a. Retry the machine code update. If you created a DVD for this procedure, use a new media.

   b. If the problem persists, contact your next level of support.

Upgrading your HMC software

Learn how to upgrade the software on an HMC from one release to the next while you maintain your HMC configuration data.

About this task

To upgrade the machine code on an HMC, complete all steps.

Note: For HMC models 7063-CR1 and 7063-CR2, you can connect an external USB DVD drive.

Step 1. Obtain the upgrade

About this task

You can order the HMC machine code upgrade through the Fix Central website.

To obtain the upgrade through the Fix Central website, complete the following steps:

Procedure

1. From a computer or server with an internet connection, go to the Hardware Management Console website at http://www-933.ibm.com/support/fixcentral/.
2. Click **Continue**.

   The Hardware Management Console site is displayed.
3. Navigate to the HMC version you want to upgrade to.
4. Locate the download and ordering section.

   Note: If you do not have access to the internet, contact IBM service and support to order the upgrade on DVD.
5. Follow the on-screen prompts to submit your order.
6. After you have the upgrade, continue with “Step 2. View the existing HMC machine code level” on page 118.

Step 2. View the existing HMC machine code level

About this task

To determine the existing level of machine code on an HMC, follow these steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Console Management**. In the navigation area, click **Updates**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver information heading, including the HMC version, release, maintenance level, build level, and base versions.
4. Continue with “Step 3. Back up the managed system's profile data” on page 118.

Step 3. Back up the managed system's profile data

About this task

To back up the managed system's profile data, complete the following steps:

Procedure

1. Select the system that you want to save the profile data.
2. Click **Actions** > **View All Actions** > **Legacy** > **Manage Partition Data** > **Backup**.
3. Type a backup file name and record this information.
4. Click **OK**.
5. Repeat these steps for each system.
6. Continue with "Step 4. Back up HMC data" on page 119.

Step 4. Back up HMC data

About this task

Back up HMC data before you install a new version of HMC software so that previous levels can be restored in the event of a problem while you upgrade the software. Do not use this critical console data after a successful upgrade to a new version of the HMC software.

Note: To back up to removable media, you need to have that media available.

To back up HMC data, complete the following steps:

Procedure

1. If you plan to back up to media, perform the following steps to format the media:
   a. Insert the media into the drive.
   b. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
   c. In the content pane, click **Format Media**.
   d. Select the media type.
   e. Select the format type.
   f. Click **OK**.
2. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.
3. In the content pane, click **Backup Management Console Data**.

   The **Backup Management Console Data** window opens.
4. Select an archive option.

   You can back up to media on a local system, a remote system that is mounted to the HMC file system (for example, NFS), or send the backup to a remote site by using File Transfer Protocol (FTP).

   - To back up to a local system, choose **Back up to media on local system** and follow the instructions.
   - To back up to a mounted remote system, choose **Back up to mounted remote system** and follow the instructions.
   - To back up to a remote FTP site, choose **Send back up critical data to remote site** and follow the instructions.
5. Continue with "Step 5. Record the current HMC configuration information" on page 119.

Step 5. Record the current HMC configuration information

About this task

Before you upgrade to a new version of HMC software, as a precautionary measure, record HMC configuration information.

To record the current HMC configuration, complete the following steps:

Procedure

1. Select a managed system or any partitions that you want to record HMC configuration information.
2. From the menu pod, select **Actions** > **Schedule Operations**.
   All scheduled operations for the target that you selected are displayed.
3. Select **Sort** > **By Object**.
4. Select each object and record the following details:
   - Object Name
   - Schedule date
   - Operation Time (displayed in 24-hour format)
   - Repetitive (if Yes, complete the following steps):
     a. Select **View** > **Schedule Details**.
     b. Record the interval information.
     c. Close the scheduled operations window.
     d. Repeat for each scheduled operation.
5. Close the **Customize Scheduled Operations** window.
6. Continue with "Step 6. Record remote command status" on page 120.

Step 6. Record remote command status

About this task

To record remote command status, complete the following steps:

Procedure

1. In the navigation area, click the **Users and Security** icon 🔓, and then select **Systems and Console Security**.
2. In the content pane, click **Enable Remote Command Execution**.
3. Record whether the **Enable remote command execution using the ssh facility** check box is selected.
4. Click **Cancel**.
5. Continue with "Step 7. Save upgrade data" on page 120.

Step 7. Save upgrade data

About this task

You can save the current HMC configuration in a designated disk partition on the HMC or to local media. Save upgrade data only immediately before you upgrade your HMC software to a new release. You can restore the HMC configuration settings after you upgrade.

Note: Only one level of backup data is allowed. Each time that you save upgrade data, the previous level is overwritten.

To save upgrade data, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Management**.
2. In the content pane, click **Save Upgrade Data**. The **Save Upgrade Data** wizard opens.
3. Select the media on which you want to save the upgrade data. If you choose to save to removable media, insert the media now. Click **Next**.
4. Click **Finish**.
5. Wait for the task to complete.

   If the Save Upgrade Data task fails, contact your next level of support before proceeding.

   Note: If the save upgrade data task fails, do not continue the upgrade process.
6. Click **OK**.
7. Continue with "Step 8. Upgrade the HMC software" on page 121.

Step 8. Upgrade the HMC software

About this task

To upgrade the HMC software, restart the system with the removable media in the DVD drive.

Procedure

1. Insert the HMC Product Installation media into the DVD drive.
2. In the navigation area, click the **HMC Management** icon, and then select **Console Management**.
3. In the content pane, select **Shutdown or Restart the Managment Console**.
4. Ensure **Restart the HMC** is selected.
5. Click **OK**.

   The HMC restarts and system information scrolls on the window.
6. Select **Upgrade** and click **Next**.
7. Choose from the following options:
   - If you saved the upgrade data during the previous task, continue with the next step.
   - If you did not save the upgrade data previously in this procedure, you must save the upgrade data now before you continue.
8. Select **Upgrade from media** and click **Next**.
9. Confirm the settings and click **Finish**.
10. Follow the prompts.

    Note:

    - If the screen goes blank, press the space bar to view the information.
    - The first DVD can take approximately 20 minutes to install.
11. At the login prompt, log in using your user ID and password.

    The HMC code installation is complete.
12. Continue with "Step 9. Verify that the HMC machine code upgrade installed successfully" on page 122.

Step 9. Verify that the HMC machine code upgrade installed successfully

About this task

To verify that the HMC upgrade is installed successfully, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver information heading, including the HMC version, release, maintenance level, build level, and base versions.
4. Verify that the version and release match the update that you installed.
5. If the level of code that is displayed is not the level that you installed, retry the upgrade task by using a new DVD. If the problem persists, contact your next level of support.

Upgrading HMC from remote location by using network upgrade images

Learn how to upgrade the software on an HMC from a remote location by using network upgrade images.

About this task

Learn how to upgrade the software on an HMC from a remote location by using network upgrade images.

Procedure

1. From a computer or server with an internet connection, go to the Hardware Management Console Support and downloads website (http://www14.software.ibm.com/webapp/set2/sas/f/hmcl/home.html).
2. Download the appropriate HMC V9 network images and save them on an FTP server.

   You cannot download these files directly to the HMC. You must download the image files to a server that accepts FTP requests.
3. Ensure that you download the following files:
   - img2a
   - img3a
   - base.img
   - disk1.img
   - hmcnetworkfiles.sum
4. Save the upgrade data on the HMC. Run the following commands to save the upgrade data:
   - To save data on both DVD and HDD, run the following commands:

     ```
     mount /media/cdrom
     ```

     ```
     saveupgdata -r diskdvd
     ```
   - To save data on the HDD, run the following command:

     ```
     saveupgdata -r disk
     ```
5. Copy the upgrade files to the bootable disk partition on the HMC. Run the **`getupgfiles`** command to copy the files.

   Example: **`getupgfiles -h <ftp server> -u <user id> -d <remote directory>`**

Where,

- **`ftp server`** is the host name or IP address of the FTP server where you download the HMC network images.
- **`user id`** is a valid user ID on the FTP server. If you do not specify the password with the `--passwd` argument, you are prompted for a password.
- **`remote directory`** is the directory on your FTP server where the HMC network images are saved.

6. Restart the HMC to upgrade the code that is copied to the bootable disk partition. Run the **`chhmc -c altdiskboot -s enable --mode upgrade`** to restart the HMC.
7. Restart the HMC and start the upgrade. Run the **`hmcshutdown -r -t now`** command to start the upgrade.

Securing the HMC

Learn how to enhance the security of your Hardware Management Console (HMC) that is based on your corporate security standards.

The default configuration of the HMC provides ample security for most enterprise users. With the Hardware Management Console (HMC) Version 8.4.0, or later, you can further enhance the security of the HMC that is based on your corporate security standards. To enhance the security for the HMC, you must set the HMC to a minimum of Level 1 security. You may choose Level 2 and Level 3 security depending on your environment and the corporate security requirements.

Note: Before changing the security level, check with your corporate security compliance team.

Level 1 security

To secure the HMC (level 1 security), complete the following steps:

1. Change the predefined password for the default `hscroot` user. For more information about password policy, see "Enhanced password policy" on page 125.
2. If the HMC does not belong to a physically secure environment, set the `grub` password by running the following command: `chhmc -c grubpasswd -s enable --passwd <new grub password>`
3. If you have configured the Integrated Management Module (IMM) on the HMC, set a strong IMM password.
4. Set a strong password for the *admin* user and general users on all servers.
5. Update the HMC with the latest released security fixes. For more information about the security fixes, see IBM Fix Central.

Level 2 security

If you have multiple users, complete the following steps to enhance the security for the HMC:

1. Create an account for each user on the HMC and assign the required roles and resources to users. For more information about the various roles in the HMC, see HMC tasks, user roles, IDs, and associated commands.

   Note: Ensure that you assign only the required resources and roles for users that are created on HMC. If necessary, you can also create custom roles.
2. Enable user data replication between different Hardware Management Consoles. The user data replication can be performed in Master-slave mode or Peer-Peer mode. For more information about user data replication, see Manage Data Replication.
3. Import a certificate that is signed by the Certificate Authority.

Level 3 security

If you have multiple Hardware Management Consoles and system administrators, complete the following steps to enhance the security for the HMC:

1. Use centralized authentication such as Lightweight Directory Access Protocol (LDAP) or Kerberos. For more information about configuring LDAP, see How to Configure LDAP on HMC.
2. Enable user data replication between different Hardware Management Consoles.
3. Ensure that the HMC is in NIST SP 800-131A mode so that the HMC uses only strong ciphers.
4. Block ports that are not required in the firewall. For information about the HMC ports that can be used, see the following table:

Table 36. Port used by the user for interaction with HMC

Port	Description	Type	Protocol version (Default mode)	Protocol Version (NIST mode)
22	Open SSH	TCP	SSH v3	SSH v3
123	NTP	UDP	NTP	NTP
161	SNMP Agent	UDP	SNMP v3	SNMP v3
162	SNMP Trap	UDP	SNMP v3	SNMP v3
427	SLP	UDP	N/A	N/A
443	HMC GUI and REST API	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
657	RMC	TCP/UDP	RSCT (Plain text + hash and sign)	RSCT (Plain text + hash and sign)
2300	5250 Terminal for IBM i	TCP	Plain text	Plain text
2301	5250 Secure terminal for IBM i	TCP	TLS 1.2	TLS 1.2
5989	CIM (legacy port, non-functional)	TCP	Non-functional	Non-functional
9900	FCS: HMC-HMC discovery	UDP	FCS	FCS
9920	FCS: HMC-HMC communication	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
9960	VTerm applet in GUI	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
12443	HMC REST API (legacy port)	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
12347	RSCT Peer Domain	UDP	RSCT (Plain text + hash and sign)	RSCT (Plain text + hash and sign)
12348	RSCT Peer Domain	UDP	RSCT (Plain text + hash and sign)	RSCT (Plain text + hash and sign)

Notes:

- You must use only SSH (port 22), HTTPS (port 443 and port 12443), 5250 secure terminal for IBM i (port 2301), and VTerm (port 9960) that are exposed to an intranet. All other ports must be used in a private or isolated network. You can use a separate Ethernet port and VLAN for the Resource

Monitoring and Control (RMC) (port 657), FCS (port 9900 and port 9920), and RSCT Peer Domain (port 12347 and port 12348).

- Ports listed in the **`netstat`** command are used for internal processes only.

Enhanced password policy

You can enforce password requirements for locally authenticated users by using the Hardware Management Console (HMC). The enhanced password policy function allows the system administrator to set password restrictions. The enhanced password policy applies to the systems in which an HMC is installed.

System administrators can use the enhanced password policy to define a single password policy for all users. The HMC provides a medium security password policy, which can be activated by the system administrators to set password restrictions. The system administrator can also choose to activate the medium security policy or a new user-defined policy. The HMC medium security password policy cannot be removed from the system. The following table lists the attributes of the medium security policy and the default values.

Table 37. Password attributes for the HMC medium security password policy

Attribute	Description	Default value
`min_pwage`	The minimum number of days for which a password must remain active.	1
`pwage`	The maximum number of days for which a password might remain active.	180
`min_length`	The minimum length of a password.	8
`hist_size`	The number of previously saved passwords that cannot be reused.	10
`warn_pwage`	When the password is about to expire, the number of days before which a user is warned that the password is about to expire.	7
`min_digits`	The minimum number of digits that are required to be used in the password.	None
`min_uppercase`	The minimum number of upper case characters.	1
`min_lowercase`	The minimum number of lower case characters.	6
`min_special_chars`	The minimum number of special characters that must be used in the password.	None

Consider the following items about the HMC medium security password policy:

- The policy features for password age and login disablement are not applicable to the **hscroot**, **hscpe**, and **root** user IDs. The policy feature for password character validation is applicable for these user IDs.
- The policy affects only the locally authenticated users that are managed by the HMC and the policy cannot be enforced on LDAP or Kerberos users.
- The HMC medium security password policy or the user-defined policy allows the system administrators to set password reuse restrictions.
- The HMC medium security password is read-only and the attributes of HMC medium security password cannot be changed. You can create a new user-defined password to set password restriction.

You can use the following commands to configure the HMC medium security password policy:

mkpwdpolicy
Imports the password policy from a file, which contains all the parameters, or creates a password policy.

lspwdpolicy
Lists all the available password policy profiles and searches for specific parameters. You can also view the password policy that is currently active.

rmpwdpolicy
Removes an existing inactive password policy.

Note: You cannot remove an active medium security policy and the default read-only password policy.

chpwdpolicy
Changes parameters of an inactive password policy.

Solving common problems while securing the HMC

Learn how to solve some problems that you might encounter when you secure the HMC.

How to secure the connection between the Hardware Management Console (HMC) and the system?

The HMC connects to the system through the Flexible Service Processor (FSP). A proprietary binary protocol called Network Client protocol (NETC) is used for managing both FSP and Power hypervisor. The following table lists ports that are used by the HMC:

Table 38. Ports on FSP that are used to interact with the HMC

Port on FSP	Description	Protocol version (Default mode)	Protocol Version (NIST mode)
443	Advanced System Management Interface	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
30000	NETC	NETC (TLS 1.2). Falls back to SSLv3 for support of older firmware.	NETC (TLS 1.2)
30001	VTerm	NETC (TLS 1.2). Falls back to SSLv3 for support of older firmware.	NETC (TLS 1.2)

How to lock the HMC?

If you want to enhance the security for your infrastructure, you can use an Intrusion Prevention System (IPS) device or add all Hardware Management Consoles and IBM Power Systems servers behind a firewall. Also, you can disable network services on the HMC if you do not use it remotely or if you want to lock the HMC down. To disable network services on the HMC, complete the following steps:

1. Disable remote command execution by using the SSH port.
2. Disable remote virtual terminal (VTerm port).
3. Disable remote web access (HMC graphical user interface and REST API).
4. Block ports in firewall by using HMC network settings for each configured Ethernet port.

How to set the HMC in NIST SP 800-131A compliance mode?

With HMC Version 8.1.0, or later, when you set the HMC in the compliance mode, only strong ciphers listed by NIST SP 800-131A are supported. You might not be able to connect to older Power Systems servers such as, POWER5 servers that do not support Transport Layer Security (TLS 1.2). For more information about changing the security mode, see HMC V8R8 NIST mode.

How to view and change ciphers that are used by the HMC?

With HMC Version 8.1.0, or later, the HMC supports more secure cipher sets that are defined in NIST 800-131A. Ciphers that are used in the default mode are strong. For more information about encryption ciphers that are used by the HMC, run the **lshmcencr** command. If your corporate standards requires the use of a different set of ciphers, run the **chhmcencr** command to modify the encryption ciphers.

To list the encryption ciphers that are used by the HMC to encrypt user password, run the following command:

```
lshmcencr -c passwd -t c
```

To list the encryption ciphers that can currently be used by the HMC web user interface and REST API, run the following command:

```
lshmcencr -c webui -t c
```

To list the encryption ciphers and MAC algorithm that can currently be used by the HMC SSH interface, run the following command:

```
lshmcencr -c ssh -t c
```

```
lshmcencr -c sshmac -t c
```

How to check the strength of the certificate on the HMC?

The self-signed certificates on the HMC use SHA256 with 2048-bit RSA encryption, which is strong. If you are using CA signed certificates, ensure that you are not using the 1024-bit encryption, which is weak. The following certificates can be used for the HMC:

- The CA signed certificate can be used for the HMC graphical user interface and REST API (ports 443 and 12443).
- The port 9920 is used for HMC to HMC communication. You cannot replace this certificate with your own certificate.

How to choose between a self-signed certificate (default) or a CA signed certificate?

The HMC auto-generates a certificate during installation. However, you can generate a Certificate Signing Request (CSR) from the HMC and get a new certificate that is issued by a Certificate Authority. You can import this certificate into HMC. Ensure that you also obtain a domain name for the HMC. For more details about managing the certificates in HMC, see Manage Certificates.

How to audit the HMC?

The audit on the Hardware Management Consoles focuses on configured ciphers and the usage activity of the various HMC users. Use the following commands to view the usage activity of various HMC users:

Table 39. Ciphers that are used by the HMC

Purpose	Command
Password encryption (global setting)	`lshmcencr -c passwd -t c`
Password encryption for each user	`lshmcusr -Fname:password_encryption`
SSH ciphers	`lshmcencr -c ssh -t c`
SSH MAC	`lshmcencr -c sshmac -t c`
Cipher that are used for the HMC graphical user interface and REST API	`lshmcencr -c webui -t c`

Use the following commands to monitor various console and serviceable events information for uses in the HMC:

Table 40. Commands to view the logged on users and console or serviceable events information in the HMC

Information	Command
GUI users	`lslogon –r webui –u`
GUI tasks	`lslogon –r webui –t`
CLI users	`lslogon –r ssh –u`
CLI tasks	`lslogon –r ssh –t`
Operations on HMC	`lssvcevents -t console -d <number of days>`
Operations on System	`lssvcevents –t hardware –m <managed system> -d <number of days>`

Centralized monitoring events for the HMC: If you have many Hardware Management Consoles, set the `rsyslog` file to collect all the usage data.

How does IBM fix the HMC security vulnerabilities?

IBM has a security incidence response process named IBM Product Security Incident Response Team (PSIRT). The IBM Product Security Incident Response Team (PSIRT) is a global team that manages the receipt, investigation and internal coordination of security vulnerability information related to IBM offerings. Open Source and IBM components that are shipped with the HMC are actively monitored and analyzed. Interim fixes and security fixes are provided by IBM for all supported releases of the HMC.

How to track new interim fixes on the HMC?

The security bulletin contains information about the vulnerability and interim fixes for supported HMC versions. To track interim fixes on the HMC, you can:

- Search the latest security bulletins at IBM Security Bulletin.
- Follow @IBMPowereSupp on Twitter for notifications.
- Subscribe to email notifications at IBM Support.

Security profiles: Global Data Protection Regulation (GDPR) and Payment Card Industry Data Security Standard (PCI-DSS)

Learn about how the Hardware Management Console (HMC) handles the privacy information of the users.

The Hardware Management Console (HMC) is a closed appliance that does not store any cardholder data. Hence, only a subset of requirements and security assessment procedures of IT security that are defined by PCI-DSS are applicable for the HMC. Only trusted code that is distributed by IBM can be installed on the HMC. When any vulnerability is known through the IBM PSIRT process, interim fixes are published. The requirements and recommendations include the following items:

GDPR queries

Table 41. GDPR queries . The table provides information about the questions related to GDPR.

Questions	Answers
What kind of data is stored in the HMC?	HMC stores configuration information of Power hardware, PowerVM virtualization, and the performance metrics information.
Does the HMC process any personal data?	You can provide contact information for the call home function. Providing contact information for the call home function is optional.
Which predefined accounts are used for system administration of the HMC?	The system administrator user uses the *hscroot* username.
Do any of the accounts in the HMC relate to a specific person?	No.
Is it mandatory to provide personal data in the HMC?	No. You do not need to provide personal data information. However, providing this information is optional.
Does the HMC log file have any personal data information?	No.
Is it possible to delete personal data completely and permanently?	Yes. Unconfigure the call home function.

PCI-DSS queries

Table 42. PCI-DSS queries . The table provides information about the questions related to PCI-DSS

Questions	Answers
How to install and maintain a firewall configuration to protect the data of the cardholder?	The HMC does not store or access any cardholder data. However, the HMC has a firewall configuration and the user can control and enable specific ports.
Can I use vendor-supplied default value for system passwords and other security parameters?	Before you install a system on the network, change all the predefined passwords of the *hscroot* user.
How does the HMC protect the stored data of the cardholder?	The HMC does not store or access any cardholder data.
How does the HMC encrypt the data of the cardholder when the data is transmitted across open public networks?	The HMC does not store or access any cardholder data.
How to use and regularly update anti-virus software programs?	The HMC is a closed appliance. Therefore, malware cannot infect the HMC.
How to develop and maintain secure systems and applications?	You must install the required patches to your system manually from the IBM Fix Central website. Only trusted code that are distributed by IBM can be installed on the HMC.
Does the HMC restrict access to the cardholder data?	The HMC does not store or access any cardholder data.

Table 42. *PCI-DSS queries* . The table provides information about the questions related to PCI-DSS *(continued)*

Questions	Answers
How to assign a unique ID to each person who has access to the computer?	You can implement this requirement by ensuring that there are no shared IDs and by following the password policies.
How to restrict the physical access to the data of the cardholder?	The HMC does not store or access any cardholder data
How to track and monitor the access to network resources and to the cardholder data?	The HMC does not store or access any cardholder data.
How does the HMC test the security of the system and processes?	Scan tools are used to run security scans on all the released versions of the HMC. The scan tools include: *Qualys*, *Nessus*, *testssl*, *sslscan* and *ASoC*.
How to maintain a security policy that includes information security for employees and contractors?	System administrator disables the remote user login, activates the user login on a need basis, and deactivates the user login when the access is no longer required.

HMC port locations

You can find port locations by using location codes. Use the HMC port location illustrations to map a location code to the HMC port position on the server.

Model 5105-22E, 9008-22L, 9009-22A, 9009-22G, 9223-22H, and 9223-22S HMC port locations

Use this diagram and table to map the HMC ports on the 5105-22E, 9008-22L, 9009-22A, 9009-22G, 9223-22H, and 9223-22S.

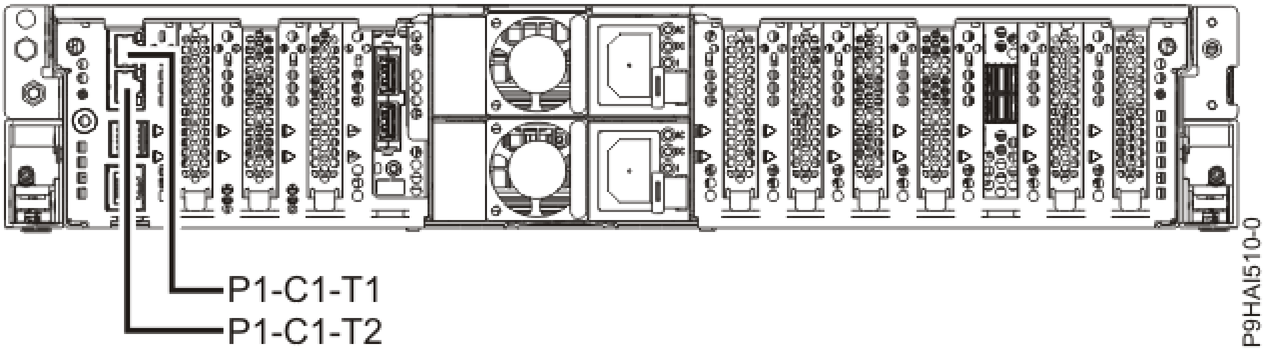

Figure 26. 5105-22E, 9008-22L, 9009-22A, 9009-22G, 9223-22H, and 9223-22S HMC port locations

Table 43. 5105-22E, 9008-22L, 9009-22A, 9009-22G, 9223-22H, and 9223-22S HMC port locations

Port	Physical location code	Identify LED
HMC port 1	Un-P1-C1-T1	No
HMC port 2	Un-P1-C1-T2	No
For more information about HMC port locations on the 5105-22E, 9008-22L, 9009-22A, 9009-22G, 9223-22H, or 9223-22S, see Part location and location codes for 9008-22L, 9009-22A, 9009-22G, 9223-22H, or 9223-22S.		

Model 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-42H, and 9223-42S HMC port locations

Use this diagram and table to map the HMC ports on the 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-42H, and 9223-42S.

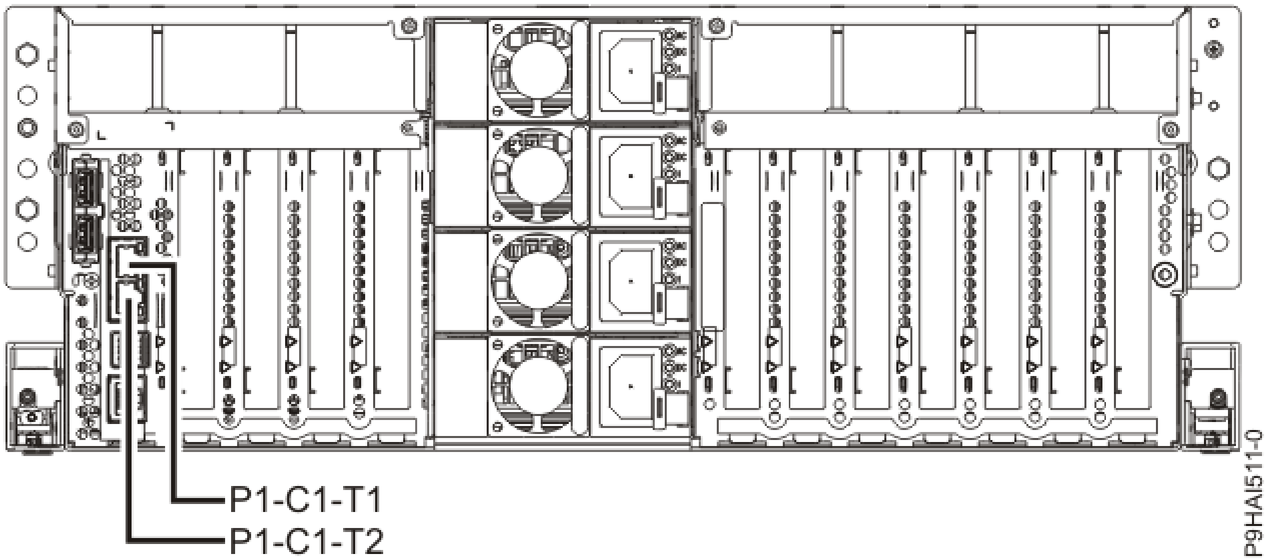

Figure 27. 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-42H, and 9223-42S HMC port locations

Table 44. 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-42H, and 9223-42S HMC port locations

Port	Physical location code	Identify LED
HMC port 1	Un-P1-C1-T1	No
HMC port 2	Un-P1-C1-T2	No
For more information about HMC port locations on the 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-42H, or 9223-42S, see Part location and location codes for 9009-41A, 9009-41G, 9009-42A, 9009-42G, 9223-42H, or 9223-42S.		

Model 9040-MR9 HMC port locations

Use this diagram and table to map the HMC ports on the 9040-MR9.

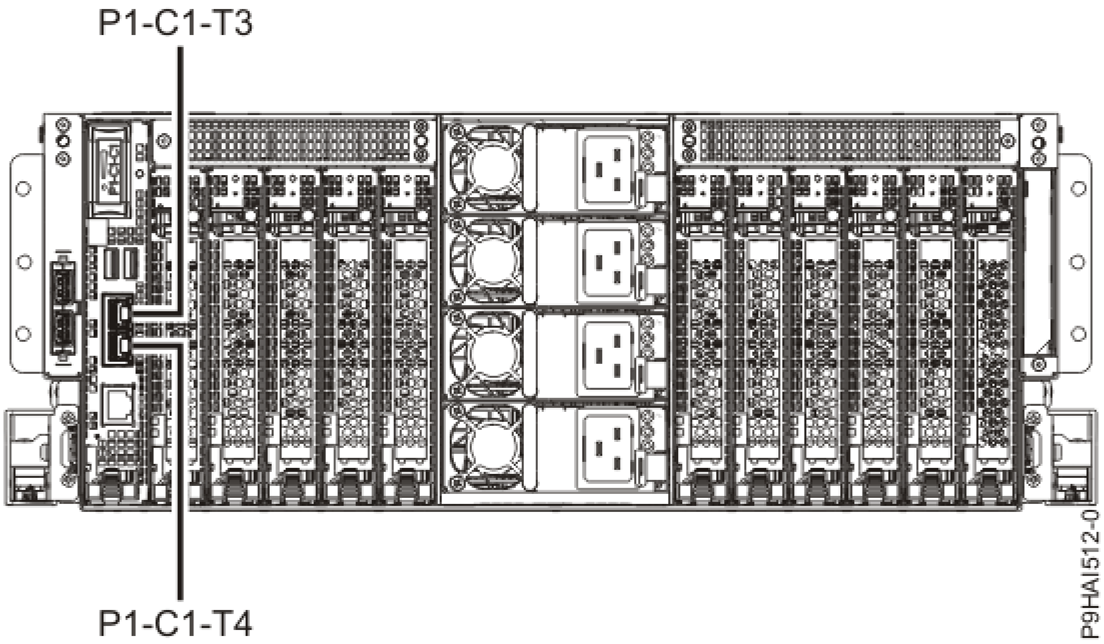

Figure 28. 9040-MR9 HMC port locations

Table 45. 9040-MR9 HMC port locations		
Port	**Physical location code**	**Identify LED**
HMC port 1	Un-P1-C1-T3	No
HMC port 2	Un-P1-C1-T4	No
For more information about HMC port locations on the 9040-MR9, see Part location and location codes.		

Model 9080-M9S HMC port locations

Use this diagram and table to map the HMC ports on the 9080-M9S.

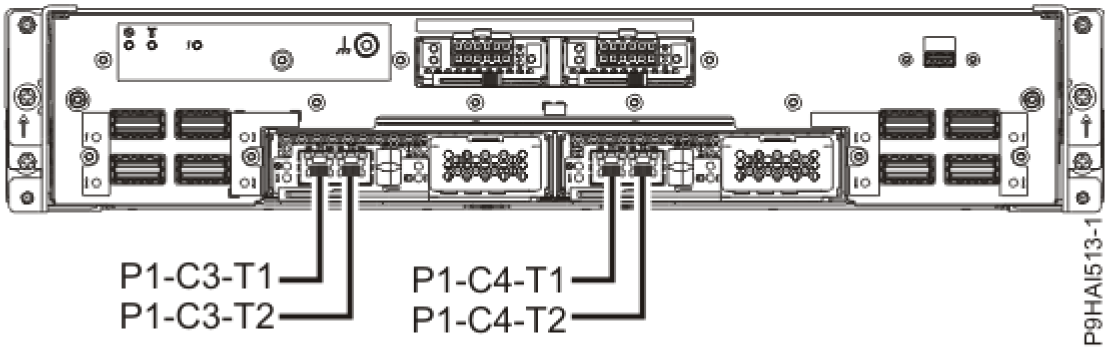

Figure 29. 9080-M9S HMC port locations

Table 46. 9080-M9S HMC port locations		
Port	**Physical port location**	**Identify LED**
Service processor card 1 - HMC port 1	Un-P1-C3-T1	No
Service processor card 1 - HMC port 2	Un-P1-C3-T2	No

Table 46. 9080-M9S HMC port locations (continued)

Port	Physical port location	Identify LED
Service processor card 2 - HMC port 1	Un-P1-C4-T1	No
Service processor card 2 - HMC port 2	Un-P1-C4-T2	No
For more information about HMC port locations on the 9080-M9S, see Part location and location codes.		

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s user name and IP address for purposes of session management. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled “Cookies, Web Beacons and Other Technologies”.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Red Hat, JBoss, OpenShift, Fedora, Hibernate, Ansible, CloudForms, RHCA, RHCE, RHCSA, Ceph, and Gluster are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Electronic emission notices

Class A Notices

The following Class A statements apply to the IBM servers that contain the POWER9 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.

New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

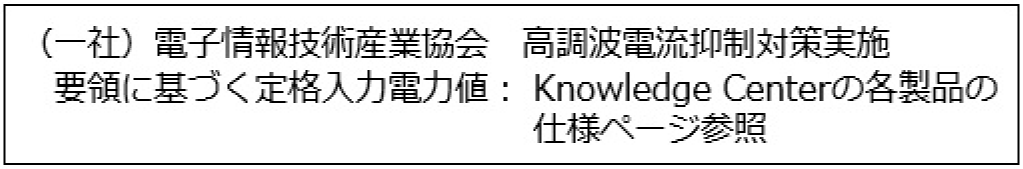

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

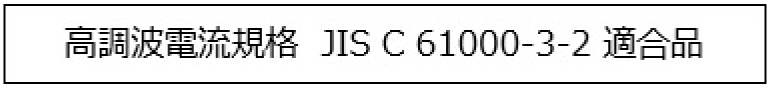

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

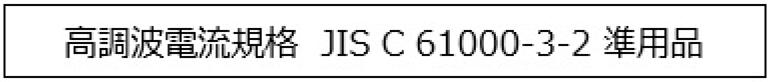

高調波電流規格 JIS C 61000-3-2 準用品

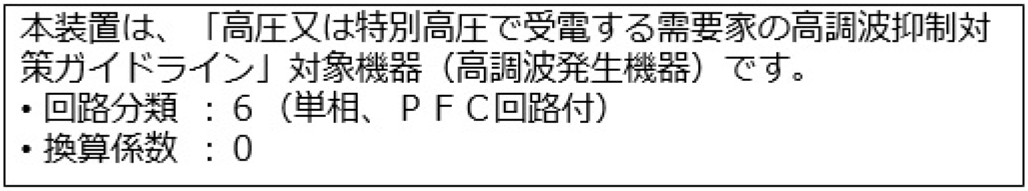

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

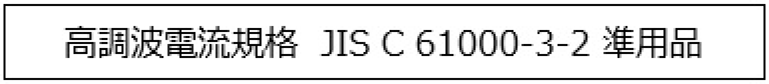

高調波電流規格 JIS C 61000-3-2 準用品

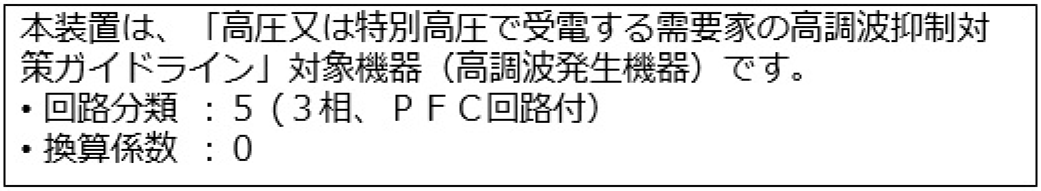

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

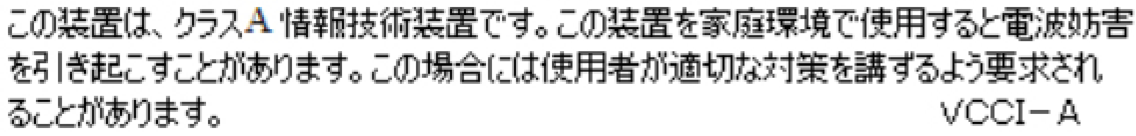

この装置は、クラスＡ 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI－A

Korea Notice

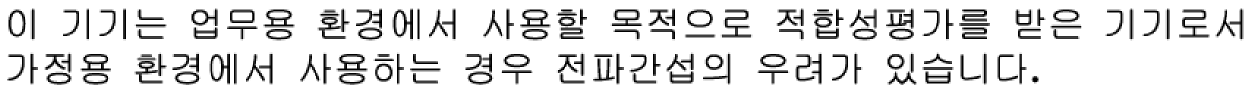

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서
가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

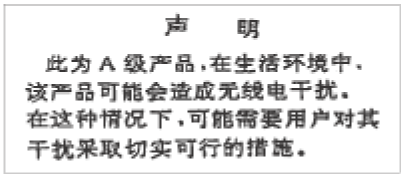

声 明

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰。
在这种情况下，可能需要用户对其
干扰采取切实可行的措施。

Russia Notice

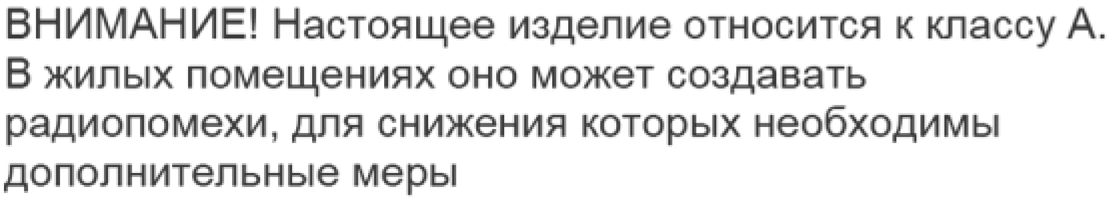

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

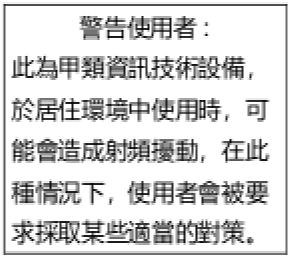

警告使用者：
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

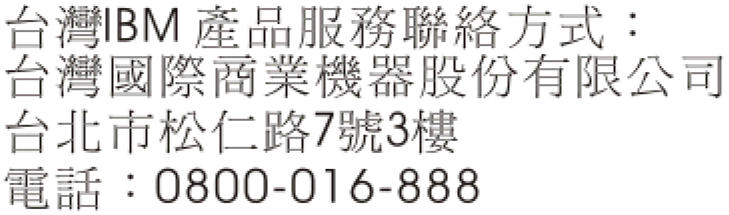

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road

Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

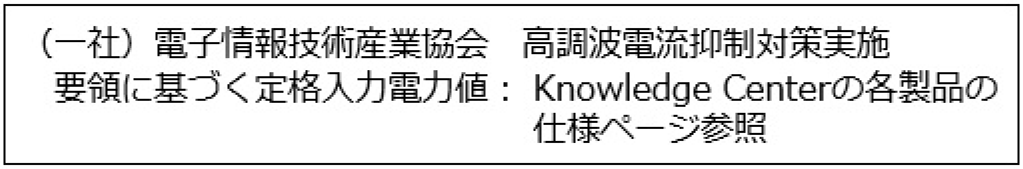

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値： Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

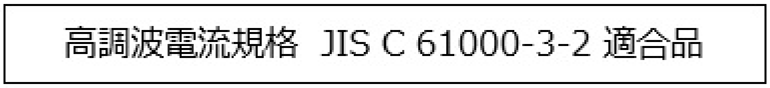

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

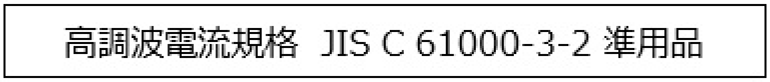

高調波電流規格 JIS C 61000-3-2 準用品

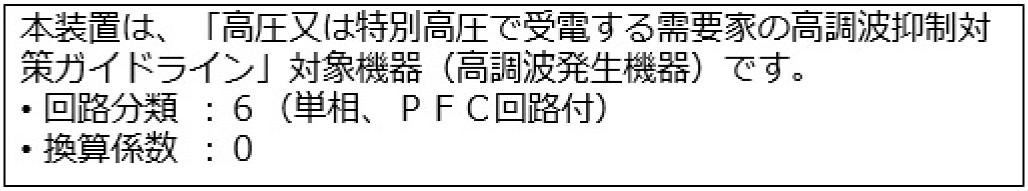

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
• 回路分類 ：6（単相、ＰＦＣ回路付）
• 換算係数 ：0

This statement applies to products greater than 20 A per phase, three-phase.

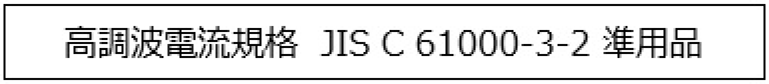

高調波電流規格 JIS C 61000-3-2 準用品

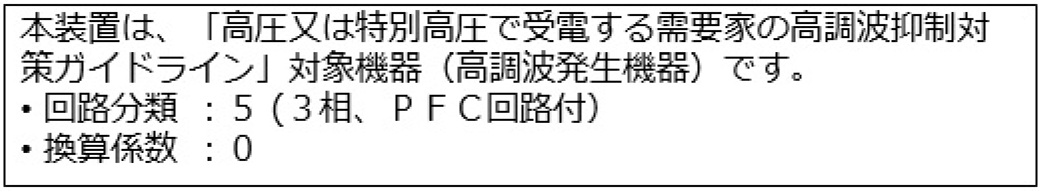

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
• 回路分類 ：5（3相、ＰＦＣ回路付）
• 換算係数 ：0

Japan Voluntary Control Council for Interference (VCCI) Notice

この装置は，クラスB情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。 VCCI－B

Taiwan Notice

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Managing the Hardware Management Console

Note

Before using this information and the product it supports, read the information in "Notices" on page 117.

This edition applies to IBM Hardware Management Console Version 9 Release 2 Maintenance Level 950 and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Managing the HMC

Learn how to use the Hardware Management Console (HMC).

About this task

Learn about the tasks that you can use on the console and how to navigate by using the web-based user interface with graphical views of managed systems and simplified navigation.

Note: Many of the HMC tasks that are listed here can also be performed by using PowerVC. For more information about the tasks that you can perform by using PowerVC, see HMC and PowerVC.

What's new in Managing the HMC

Read about new or significantly changed information in Managing the HMC since the previous update of this topic collection.

November 2020

- Added the following topics:
  - “Validate Maintenance Readiness” on page 61
  - “Manage Virtual I/O Server Backups” on page 84
- Updated the following topics:
  - “Create Partition” on page 42
  - “Serviceable Events Manager” on page 47
  - “Add FRU” on page 50
  - “Exchange FRU” on page 51
  - “Remove FRU” on page 51
  - “Partition content pane” on page 56
  - “Partition Properties” on page 57
  - “Serviceable Events Manager” on page 63

May 2020

- Updated the following topics:
  - “HMC tasks, user roles, IDs, and associated commands” on page 9
  - “Partition Templates” on page 83
  - “Hardware Virtualized I/O” on page 66

February 2020

Updated the following topic:

- “VIOS Images” on page 83

November 2019

- Added the following topic:
  - “Logging on to the Hardware Management Console (HMC) when PowerSC MFA is configured on the HMC” on page 3

- Updated the following topics:
  - “Introduction to the HMC” on page 2
  - “Predefined user IDs and passwords” on page 4
  - “Partition Properties” on page 57
  - “Hardware Virtualized I/O” on page 66
  - “Console Default Settings” on page 77
  - “Backup Management Console Data” on page 80
  - “Partition Templates” on page 83
  - “Manage Task and Resource Roles” on page 89
  - “Manage Event Notification” on page 104

May 2019

- Added the following topics:
  - “Netboot” on page 58
  - “Manage Groups” on page 72
- Updated the following topics:
  - “System content pane” on page 30
  - “Create Partition” on page 42
  - “Processor, Memory, I/O” on page 43
  - “Partition content pane” on page 56
  - “Change User Password” on page 86

August 2018

- Added the following topics:
  - “Manage MFA” on page 96
  - “Console Default Settings” on page 77
  - “Managing OpenBMC-based and BMC-based systems by using the HMC” on page 110
- Updated the “Change Network Settings” on page 75 topic.

December 2017

- Added information for HMC Version 9, Release 1, or later on IBM Power Systems servers that contain the POWER9™ processor.

Introduction to the HMC

Learn about some of the concepts and functions of the Hardware Management Console (HMC) and the user interface that is used for accessing those functions.

You can configure and manage servers on the HMC. One HMC can manage multiple servers, and dual HMCs can provide redundant support by managing the same system. To ensure consistent function, each HMC is shipped preinstalled with the HMC Licensed Machine Code Version 9, Release 1.

To provide flexibility and availability, you can implement HMCs in several configurations.

HMC as the DHCP server

An HMC that is connected by either a private network to the systems it manages might be a DHCP server for the service processors of the systems. An HMC might also manage a system over an open network, where the managed system's service processor IP address is assigned by a customer-supplied DHCP server or manually assigned by using the Advanced System Management Interface (ASMI).

Physical proximity

Before HMC Version 7, at least one local HMC was required to be physically located near the managed systems. As an alternative to the local HMC, you can use a supported device, such as a personal computer that has connectivity and authority to operate through a remotely attached HMC. The local device must be in the same room as your server and at a distance of 8 m (26 ft) from your server. The local device must have the functional capability that is equivalent to the HMC that it replaces and that is needed by the service representative to service the system. For a virtual HMC, the functional capabilities also include the method of transferring service data, such as firmware updates or diagnostic data, and transferring the log information to and from the HMC.

Redundant or Dual HMCs

A server might be managed by either 1 or 2 Hardware Management Consoles. When two Hardware Management Consoles manage one system, they are peers, and each HMC can be used to control the managed system. The best practice is to attach one HMC to the supported networks or HMC ports of the managed systems. The networks are intended to be independent. Each HMC might be the DCHP server for a service network. Because the networks are independent, the DHCP servers must be set up to provide IP addresses on two unique and nonroutable IP ranges.

Redundant or Dual HMCs that manage the same server must not be at different version and release levels. For example, an HMC at Version 7 Release 7.1.0 and an HMC at Version 7 Release 3.5.0 cannot manage the same server. The HMCs must be at the same version and release level.

When the server is connected to the higher version of the management console, the partition configuration is upgraded to the latest version. After the partition configuration upgrade, lower levels of the management consoles will not be able to interpret the data correctly. After the server is managed by the higher version of the management console, you must first initialize the server before you can go back to the lower version of the management console. You can restore a backup that is taken at the older level or re-create the partitions. If the server is not initialized, one of the following outcomes can occur depending on the version of the lower-level HMC:

- HMC Version 7 Release 7.8.0 and later reports a connection error of **Version mismatch** with reference code **Save Area Version Mismatch**.
- HMC Version 7 Release 7.7.0 and earlier might report a server state of **Incomplete** or **Recovery**. In addition, partition configuration corruption can also occur.

Logging on to the Hardware Management Console (HMC) when PowerSC MFA is configured on the HMC

Learn how to log in to the HMC when IBM PowerSC Multi-factor Authentication (MFA) is configured on the HMC.

If IBM PowerSCMFA is enabled on the HMC and the user is configured on the PowerSC MFA server, you can choose to log in to the HMC by first entering the user ID and a policy name that is provided by your security administrator. You are then prompted to provide additional credentials.

In the HMC login page, if you click **Policy Name**, the authentication mechanism is set to the in-band authentication type. For example, if the policy that you want to use is associated with the Rivest-Shamir-Adleman (RSA) authentication method, you can enter the secure ID passcode that you received from the RSA secure ID device or the application. Then, click **Next or Sign In** to log in to the HMC.

Notes:

- If MFA is not enabled on the HMC, you can log in to the HMC with the user ID and password.
- If you obtain a cache token credential (CTC) code from the PowerSC MFA server that is configured by your security administrator, enter the CTC code in the **Password** field.

Predefined user IDs and passwords

Predefined user IDs and passwords are included with the Hardware Management Console (HMC). It is imperative to the security of your system that you change the `hscroot` predefined password immediately.

If the password expires when you try to log in to the HMC, complete the following steps:

1. Enter the **Current Password** and the **New Password**.
2. Re-enter the new password in the **Confirmation new password** field.
3. Click **OK**. If the new password complies with the current password policy, the password for the HMC is changed.

The following predefined user IDs and passwords are included with the HMC:

Table 1. Predefined HMC user IDs and passwords

User ID	Password	Purpose
`hscroot`	`abc123`	The `hscroot` user ID and password are used to log in to the HMC for the first time. They are case-sensitive and can be used only by a member of the super administrator role.

Using the web-based user interface

You can use the web-based user interface to perform tasks on the Hardware Management Console (HMC) or on your managed resources.

This user interface comprises several major components: the title bar, the navigation area, the content pane, the menu pod, and the dock pod.

The *title bar*, across the top of the workplace window, identifies the product, any user that is logged in, and help options.

The *navigation area*, in the left portion of the window, contains the primary navigation links for selecting your system and starting tasks for your HMC.

The *content pane*, in the middle portion of the window, displays information that is based on the current selection from the navigation area. For example, when **All Systems** is selected in the navigation area, all the available systems are shown in the content pane.

The *menu pod*, in the left portion of the window, is displayed after you select a system and provides quick access to commonly used HMC tasks and views of resources and properties.

The *dock pod*, in the right portion of the window, displays the *Pins* function that can be used to pin any user-selected HMC task. This function allows for quick access to these tasks.

You can resize the panes of the HMC workplace by moving the mouse pointer over the border that separates the navigation pane from the work pane until the mouse pointer changes to a double-pointed arrow. When the pointer changes shape, press and hold the left mouse button while you drag the mouse pointer to the left or right. Release the button and your navigation pane or work pane is now larger or smaller in size. You can also do this task within the work pane border that separates the resources table from the taskpad.

You can change the layout of the *content pane* according to your preference by clicking **Display Gallery View**, **Display Table View**, or **Display Relationship View**.

You can reposition columns in tables by selecting and dragging a column to a new position. You can also select which columns to display by clicking the drop-down menu that is located next to the last column of each table. You can save your preferences by clicking the **Save User Preferences** icon.

You can change how many rows are displayed in tables on each page by clicking one of the **Items per page** (**10**, **20**, **30**, or **50**) icons that are located below each table.

Note: Pop-up windows must be enabled for full functionality of the HMC.

Overview of menu options

Learn about the menu options and associated tasks that are available in the Hardware Management Console (HMC).

The menu options and tasks that are described in this section are available in the HMC interface.

Table 2. HMC menu options

Menu	Submenu	Options/Tasks
Resources	All Systems	View All Systems
	All Partitions	View All Partitions
	All Virtual I/O Servers	View All Virtual I/O Servers
	All Frames	View All Frames
	All Power Enterprise Pools	View All Power Enterprise Pools
	All Shared Storage Pool Clusters	View All Shared Storage Pool Clusters
	All Groups	View All Groups

Table 2. HMC menu options (continued)

Menu	Submenu	Options/Tasks
HMC Management	Console Settings	Launch Guided Setup Wizard
		View Network Topology
		Test Network Connectivity
		Change Network Settings
		Change Performance Management Settings
		Change Date and Time
		Change Language and Locale
	Console Management	Shut Down or Restart the Management Console
		Schedule Operations
		View Licences
		Update the Hardware Management Console
		Manage Install Resources
		Manage Virtual I/O Server Image Repository
		Format Media
		Backup Management Console Data
		Restore Management Console Data
		Save Upgrade Data
		Manage Data Replication
	Template Library	System and Partition Library
	Updates	Not available (use the Update the Hardware Management Console option instead)

Table 2. HMC menu options (continued)

Menu	Submenu	Options/Tasks
Users and Security	Users and Roles	Change User Password
		Manage User Profiles and Access
		Manage Users and Tasks
		Manage Task and Resource Roles
	Systems and Console Security	Manage Certificates
		Manage LDAP
		Manage KDC
		Enable Remote Command Execution
		Enable Remote Operation
		Enable Remote Virtual Terminal

Table 2. HMC menu options (continued)

Menu	Submenu	Options/Tasks
Serviceability	Console Events Logs	View Console Events window
	Serviceable Events Manager	Serviceable Events Manager window
	Events Manager for Call Home	Events Manager for Call Home window
	Service Management	Create Serviceable Event
		Manage Remote Connections
		Manage Remote Support Requests
		Manage Dumps
		Transmit Service Information
		Schedule Service Information
		Format Media
		Perform Management Console Trace
		View Management Console Logs
		View Component Logs
		Electronic Service Agent Setup Wizard
		Authorize User
		Enable Electronic Service Agent
		Manage Outbound Connectivity
		Manage Inbound Connectivity
		Manage Customer Information
		Manage Serviceable Event Notification
		Manage Connection Monitoring

Tasks and roles

Each HMC user can be a member of a different role. Each of these roles allows the user to access different parts of the HMC and complete different tasks on the managed system. HMC roles are either predefined or customized.

The roles that are discussed refer to HMC users; operating systems that are running on logical partitions have their own set of users and roles. When you create an HMC user, you must assign that user a task role. Each task role allows the user different levels of access to tasks available on the HMC interface. For more information about the tasks each HMC user role can perform, see "HMC tasks, user roles, IDs, and associated commands" on page 9.

You can assign managed systems and logical partitions to individual HMC users. This action allows you to create a user that has access to managed system A but not to managed system B. Each grouping of managed resource access is called a managed resource role.

The **predefined** HMC roles, which are the default on the HMC, are as follows:

Table 3. Predefined HMC Roles

Role	Description	HMC User ID
Operator	The operator is responsible for daily system operation.	**hmcoperator**
Super administrator	The super administrator acts as the root user, or manager, of the HMC system. The super administrator has unrestricted authority to access and modify most of the HMC system.	**hmcsuperadmin**
Product engineer	A product engineer helps support situations, but cannot access HMC user management functions. To provide support access for your system, you must create and administer user IDs with the product engineer role.	
Service representative	A service representative is an employee who is at your location to install, configure, or repair the system.	**hmcservicerep**
Viewer	A viewer can view HMC information, but cannot change any configuration information.	**hmcviewer**
Client live update	The client live update role is intended for use when you are using the AIX Live Update capability on a partition of a managed system. A client live update user has authority that is limited to what is necessary to complete a live update on AIX.	**hmcclientliveupdate**

You can create **customized** HMC roles by modifying predefined HMC roles. Creating customized HMC roles is useful for restricting or granting specific task privileges to a certain user.

HMC tasks, user roles, IDs, and associated commands

The roles discussed in this section refer to HMC users; operating systems running on logical partitions has its own set of users and roles.

Each HMC user has an associated task role and a resource role. The task role defines the operations the user can perform. The resource role defines the systems and partitions for performing the tasks. The users may share task or resource roles. The HMC is installed with five predefined task roles. The single predefined resource role allows access to all resources. The operator can add customized task roles, customized resource roles, and customized user IDs.

Some tasks have an associated command. For more information about accessing the HMC command line, see “Using the HMC remote command line” on page 108.

Some tasks can only be performed using the command line. For a listing of those tasks, see Table 9 on page 26.

For more information about where to find task information, see the following table:

Table 4. HMC task groupings

HMC tasks and the corresponding user roles, IDs, and commands	Associated table
HMC Management	Table 5 on page 10
Service Management	Table 6 on page 13
Systems Management	Table 7 on page 15
Control Panel Functions	Table 8 on page 24

This table describes the HMC management tasks, commands, and default user roles associated with each HMC Management task.

Table 5. HMC Management tasks, commands, and default user roles

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Backup Management Console Data" on page 80 bkconsdata	X	X		X
Backup Profile Data bkprofdata	X	X		X
Change BMC Certificates chbmccert	X	X		X
Certificate Management chhmccert lshmccert mkhmccert		X		
"Change Date and Time" on page 76 chhmc lshmc	X	X		X
"Change Language and Locale" on page 77 chhmc lshmc	X	X	X	X
Change HMC Configuration chipsec chpsm chusrtca	X	X		X

Table 5. HMC Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Change Network Settings" on page 75 chhmc lshmc	X	X		X
Change Proxy Configuration chproxy		X		X
"Change User Password" on page 86 chhmcusr	X	X	X	X
List BMC Certificates lsbmccert	X	X	X	X
List HMC Configuration lsipsec lspsm lsusrtca	X	X	X	X
List HMC Encryption Task lshmcencr	X	X	X	
List System Plan lssysplan		X		
List Proxy Configuration lsproxy	X	X	X	X
"Manage KDC" on page 92 chhmc lshmc getfile rmfile		X		
"Manage LDAP" on page 92 lshmcldap chhmcldap		X		
"Launch Guided Setup Wizard" on page 73		X		

Table 5. HMC Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Launch Remote Hardware Management Console	X	X	X	X
Lock HMC Screen	X	X	X	X
Logoff or Disconnect	X	X	X	X
"Manage Certificates" on page 90		X		
"Manage Data Replication" on page 81	X	X		
"Manage Task and Resource Roles" on page 89 chaccfg lsaccfg mkaccfg rmaccfg		X		
"Manage User Profiles and Access" on page 86 chhmcusr lshmcusr mkhmcusr rmhmcusr		X		
"Manage Users and Tasks" on page 89 lslogon termtask	X	X	X	X
Open 5250 Console	X	X		X
"Enable Remote Command Execution" on page 97 chhmc lshmc	X	X		X
"Enable Remote Operation" on page 97 chhmc lshmc	X	X	X	X

Table 5. HMC Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Enable Remote Virtual Terminal" on page 97	X	X		X
"Restore Management Console Data" on page 81	X	X		X
"Save Upgrade Data" on page 81 saveupgdata	X	X		X
"Schedule Operations" on page 78	X	X		
"Shut Down or Restart" on page 78 hmcshutdown	X	X		X
"Serviceable Events Manager" on page 47 lssvcevents	X	X		X
"View Licenses" on page 79	X	X	X	X

This table describes the Service Management tasks, commands, and default user roles.

Table 6. Service Management tasks, commands, and default user roles

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Create Serviceable Event" on page 48		X		X
"Serviceable Events Manager" on page 98 chsvcevent cpfile lssvcevents mksvcevent updpmh		X		X

Table 6. Service Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Format Media" on page 80 formatmedia	X	X		X
"Manage Dumps" on page 99 dump cpdump getdump lsdump startdump lsfru	X	X		X
"Transmit Service Information" on page 100 chsacfg lssacfg	X	X		
"Enable Electronic Service Agent" on page 102	X	X		X
"Manage Outbound Connectivity" on page 102	X	X		X
"Manage Inbound Connectivity" on page 103	X	X		X
"Manage Customer Information" on page 103	X	X		X
"Authorize User" on page 101		X		
"Manage Event Notification" on page 104 chsacfg lssacfg	X	X		X
"Manage Connection Monitoring" on page 104	X	X	X	X
"Electronic Service Agent Setup Wizard" on page 101		X		X

This table describes the Systems Management tasks, commands, and default user roles.

Table 7. Systems Management tasks, commands, and default user roles

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmi n)	Viewer (hmcview er)	Service Representative (hmcservicerep)
"General Settings" on page 43 lshwres	X	X	X	X
lsled	X	X	X	X
lslparmigr	X	X	X	X
lssyscfg	X	X	X	X
chhwres	X	X	X	X
chsyscfg	X	X	X	X
migrlpar	X	X	X	X
optmem	X	X		X
lsmemopt	X	X	X	X
lsrrstartlpar	X	X		
Update Password chsyspwd		X		
Change Default User Interface Settings	X	X	X	X
List CEC Property lscomgmt lsiotopo	X	X	X	X
List Utilization Data lslparutil	X	X	X	X
Operations				
"Power Off" on page 31 chsysstate	X	X		X
"Activate" on page 58 chsysstate	X	X		X
"Save Current Configuration" on page 63 chsysstate	X	X		X

Table 7. Systems Management tasks, commands, and default user roles (continued)				
HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Restart" on page 58 chsysstate	X	X		X
"Shut Down" on page 59 chsysstate	X	X		X
chlparstate	X	X		X
LED Status: Deactivate Attention LED "Attention LED" on page 35 chled	X	X		
LED Status: Identify LED "Attention LED" on page 35	X	X	X	X
LED Status: Test LED "Attention LED" on page 35	X	X	X	X
"Schedule Operations" on page 33	X	X		
"Launch ASM Interface" on page 34 asmmenu	X	X		X
"Rebuild" on page 35 chsysstate	X	X		
"Power Management" on page 32 chpwrmgmt lspwrmgmt		X		
"Delete" on page 59 rmsyscfg	X	X		X
"Mobility" on page 61 lslparmigr migrlpar	X	X		X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Manage Profiles" on page 62 chsyscfg lssyscfg mksyscfg rmsyscfg chsysstate	X	X		X
Manage System Plan cpsysplan rmsysplan		X		
Make System Plan mksysplan		X		
Deploy System Plan deploysysplan		X		
Change N_Port Login chnportlogin	X	X		X
RR Start LPAR lsrrstartlpar rrstartlpar	X	X		
Migrate LPAR migrdbg refdev	X	X		
Make Profile Data mkprofdata	X	X		
Restore Profile Data migrcfg	X	X		
Remove Profile Data rmprofdata	X	X		

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmi n)**	**Viewer (hmcview er)**	**Service Representative (hmcservicerep)**
Manage Pmem CEC Config: Initialize Profile Data: Restore Profile Data rstprofdata For option "--retainpmemvolume" (access only for hmcsuperadmin)	X	X		
Vios Admin Op: Virtual IO Server Command viosvrcmd For option "--admin" (access only for hmcsuperadmin and hmcoperator)	X	X		X
"Operations" on page 31	X	X	X	X
Configuration				
"Create Partition from Template" on page 37		X		
"Deploy System from Template" on page 37		X		
"Capture Configuration as Template" on page 37		X		
Change CEC Property chprimhmc	X	X		
Change Trusted System Key chtskey		X		
"Create Partition" on page 42		X		
List LPAR Property lsmigrdbg	X	X	X	X
Hibernate LPAR lsrsdevsize	X	X		
List N_Port Login lsnportlogin	X	X		X
LS Profile Space lsprofspace	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
List Trusted System Key lstskey	X	X	X	X
"Manage Custom Groups" on page 63	X	X		X
"Manage Profiles" on page 62 chsyscfg chsysstate lssyscfg mksyscfg rmsyscfg	X	X	X	X
Manage License Keys chlickey	X	X		
Manage Utilization Data chlparutil	X	X		X
Save Current Configuration "Save Current Configuration" on page 63 mksyscfg	X	X		
ViewSPP lsmemdev	X	X	X	X
Connections				
"Service Processor Status" on page 36 lssysconn	X	X	X	X
"Reset or Remove Connections" on page 36 rmsysconn	X	X		
Add Connection mksysconn	X	X		
Open V Term mkvterm	X	X		X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
Close V Term rmvterm	X	X		X
“Disconnect Another HMC” on page 36		X		
Hardware (Information)				
“Hardware Operations” on page 50	X	X	X	X
Updates				
“Change Licensed Internal Code” on page 38 lslic updlic		X		X
“Check System Readiness” on page 38 updlic		X		X
“View System Information” on page 38 lslic		X		X
Update HMC updhmc lshmc		X		X
Serviceability				
“Serviceable Events Manager” on page 63 chsvcevent lssvcevents		X		X
Change SNMP Alerts chspsnmp	X	X		X
“Create Serviceable Event” on page 48		X		X
“Reference Code Log” on page 64 lsrefcode	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Control Panel Functions" on page 64 lssyscfg	X	X		
"Add FRU" on page 50		X		X
"Add Enclosure" on page 52		X		X
"Exchange FRU" on page 51		X		X
"Remove FRU" on page 51		X		X
"Remove Enclosure" on page 52		X		X
"Power On/Off Unit" on page 50		X		X
"Manage Dumps" on page 49 dump cpdump getdump lsdump startdump lsfru	X	X		X
"Collect VPD" on page 49	X	X	X	X
"Type, Model, Feature" on page 50		X		
"Setup FSP Failover" on page 52 chsyscfg lssyscfg		X		
"Initiate FSP Failover" on page 52 chsysstate		X		
List CEC Property lsprimhmc	X	X	X	X
Capacity on Demand (CoD)				
Enter CoD code chcod		X		

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
View History Log lscod	X	X	X	X
Change CEC Property chcomgmt	X	X		
CoD Pool Management: Change CoD chcodpool	X	X		
Change CoD mkcodpool		X		
Change VET Code chvet		X		
List CoD Information lscodpool	X	X	X	X
List VET Information lsvet	X	X	X	X
Processor: View Capacity Settings lscod	X	X	X	X
Processor CUoD: View Code Information lscod	X	X	X	X
Processor: On/Off CoD: Manage chcod		X		
Processor: On/Off CoD: View Capacity Settings lscod	X	X	X	X
Processor: On/Off CoD: View Billing Information lscod	X	X	X	X
Processor: On/Off CoD: View Code Information lscod	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Processor: Trial CoD: Stop chcod		X		
Processor: Trial CoD: View Capacity Settings lscod	X	X	X	X
Processor: Trial CoD: View Code Information lscod	X	X	X	X
Processor: Reserve CoD: Manage chcod		X		
Processor: Reserve CoD: View Capacity Settings lscod	X	X	X	X
Processor: Reserve CoD: View Code Information lscod	X	X	X	X
Processor: Reserve CoD: View Shared Processor Utilization lscod	X		X	X
PowerVM® (formerly known as Advanced POWER® Virtualization): Enter Activation Code chcod		X		
PowerVM: View History Log lscod	X	X	X	X
PowerVM: View Code Information lscod	X	X	X	X
Enterprise Enablement: Enter Activation Code chcod		X		
Enterprise Enablement: View History Log lscod	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Enterprise Enablement: View Code Information lscod	X	X	X	X
Other Advanced Functions: Enter Activation Code chcod		X		
Other Advanced Functions: View History Log lscod	X	X	X	X
Other Advanced Functions: View Code Information lscod	X	X	X	X
Processor: Manage chcod		X		
Processor: View Capacity Settings lscod	X	X	X	X
Processor: View Code Information lscod	X	X	X	X
Memory: Manage chcod		X		
Memory: View Capacity Settings lscod	X	X	X	X
Memory: View Code Information lscod	X	X	X	X

This table describes the Control Panel Functions tasks, commands, and default user roles.

Table 8. Control Panel Functions tasks, commands, and user roles

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Serviceability				

Table 8. Control Panel Functions tasks, commands, and user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
(21) Activate Dedicated Service Tools chsysstate	X	X		
(65) Disable Remote Service chsysstate	X	X		
(66) Enable Remote Service chsysstate	X	X		
(67) DIsk Unit IOP Reset / Reload chsysstate	X	X		
(68) Concurrent Maintenance Power Off Domain	X	X		
(69) Concurrent Maintenance Power On Domain	X	X		
(70) IOP Control Storage Dump chsysstate	X	X		
(71) Product Engineering Debug Tools pedbg				
(72) PE Shell Access pesh	X	X	X	X

This table describes the commands that are not associated with an HMC UI task, and defines the default user roles that can perform each command.

Table 9. Command line tasks, associated commands, and user roles

Command line tasks	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
Change which encryption is used by the HMC to encrypt the passwords of locally authenticated HMC users, or change which encryptions can be used by the HMC Web UI. chhmcencr		X		
List which encryption is used by the HMC to encrypt the passwords of locally authenticated HMC users, or list which encryptions can be used by the HMC Web UI chhmcfs	X	X	X	
Free up space in HMC file systems chhmcfs	X	X		
List HMC file system information lshmcfs	X	X	X	X
Test for removable media readiness on the HMC ckmedia	X	X		X
Obtain required files for an HMC upgrade from a remote site getupgfiles	X	X		X
Provide screen capture on the HMC hmcwin	X	X	X	X
Log SSH command usage logssh	X	X	X	X
Clear or dump partition configuration data on a managed system lpcfgop		X		

Table 9. Command line tasks, associated commands, and user roles (continued)

Command line tasks	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
List environmental information for a managed frame, or for systems contained in a managed frame lshwinfo	X	X	X	X
List which HMC owns the lock on a managed frame lslock	X	X	X	X
Force an HMC lock on a managed frame to be released rmlock		X		
List the storage media devices that are available for use on the HMC lsmediadev	X	X	X	X
Manage SSH authentication keys mkauthkeys	X	X	X	X
Monitoring HMC subsystems and system resources monhmc	X	X	X	X
Remove the utilization data collected for a managed system from the HMC rmlparutil	X	X		X
Enable users to edit a text file on the HMC in a restricted mode rnvi	X	X	X	X
Restore hardware resources after a DLPAR failure rsthwres		X		
Restore upgrade data on the HMC rstupgdata	X	X		X

Table 9. Command line tasks, associated commands, and user roles (continued)

Command line tasks	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
Transfer a file from the HMC to a remote system sendfile	X	X	X	X
chsvc	X	X		X
lssvc	X	X	X	X
chstat	X	X		X
lsstat	X	X	X	X
chpwdpolicy		X		
lspwdpolicy	X	X	X	X
mkpwdpolicy		X		
rmpwdpolicy		X		
expdata		X		

Session handling

Learn about session limitations in the Hardware Management Console (HMC).

Session limitations

The HMC does not support disconnected sessions. A session logoff and a session disconnect are both considered as a session logoff. This means that you cannot reconnect to the same session to resume your task or tasks that were initiated from a previous session. Every login through the HMC creates a new session.

1. If you initiate long running tasks from the HMC interface and then log off from the session, the long running tasks continue to run in the background. However, when you log in again, a new session is created and the task progress panels (which helps track the progress of the previous tasks) are no longer available. In this scenario, if you need to check the progress of the tasks that were initiated from a previous session, you can run the respective command line interface (CLI) commands, check the state of the managed resource, or check the console event logs.

   Note: Some examples of long running tasks include the following tasks:

   System management for servers:

   - Deploy system plan
   - Code update
   - Hardware - Prepare for hot repair or upgrade

   System management for partitions:

- DLPAR memory in large units in the order of Terabytes
- Live Partition Mobility (LPM)
- Suspend or resume

HMC management:

- Backup management console data
- Restore management console data
- Save upgrade data

2. If you fail to reauthenticate within the time that is specified in the verify timeout settings, you are automatically logged off from the current session.
3. The idle timeout user property task is not functional. The HMC interface uses the default value of **0** for the idle timeout setting. If you set a different value for this setting, it is ignored.

   Note: Session, idle, and verify timeout properties are set for a user and it can be different for different users on the same HMC.

Version mismatch state for a managed system

The **Version mismatch** state can occur when the redundant or dual Hardware Management Consoles (HMCs) that manage the same server are at different version and release levels.

The **Version mismatch** state can occur for any of the following reasons:

- FSP firmware and HMC versions are incompatible.
- An HMC Version 7.7.8 or later is connected to a server that was managed by a newer version of the HMC.
- An HMC Version 7.7.8 or later is connected to a server that was managed by a lower version of the HMC and does not have enough space present to upgrade the data to HMC Version 7.7.8 or later.
- The hypervisor or server brand or model is not supported by this version of the HMC.

To recover from the **Version mismatch** state, select the appropriate action, depending on the reference code that is displayed:

- **Save Area Version Mismatch**

  HMC Version 7.7.8 and later blocks attempts to manage a server with a configuration at a newer level by posting a new **Connection error** state and reference code. If an HMC Version 7.7.8 or later is connected to a server that was managed by a newer version of the HMC that updated the configuration format, then the HMC reports a connection error of **Version mismatch** with the reference code **Save Area Version Mismatch**. This error prevents accidental corruption of the configuration.

  If you want to continue on a lower HMC version, then you must first initialize the server in the lower version of the HMC before you proceed to run any operation.

- **Profile Data Save Area is full**

  The HMC uses a storage area on each managed server to store the server configuration, primarily PowerVM partition profiles. HMC Version 7.8.0 and later increases the usage of the storage area by adding another (mostly hidden) profile for each partition. Servers that already contain many profiles might not have sufficient space to allow the HMC Version 7.8.0 and later to run properly.

  HMC Version 7.8.0 and later checks for sufficient space in this storage area and stops the connection process with a connection state of **Version mismatch** and a reference code of **Profile Data Save Area is full** if sufficient space does not exist.

- **Connecting 0000-0000-00000000 (Unsupported Hypervisor)**

  A connection state of **Version mismatch** and a reference code of **Connecting 0000-0000-00000000 (Unsupported Hypervisor)** is returned when the server is configured for a hypervisor other than PowerVM.

To recover from this state, first start the ASM by selecting the server with the **Version mismatch** and selecting **Operations** and then **Launch Advanced System Manager (ASM)**.

On models that support multiple hypervisors, the hypervisor mode setting can be found in the ASM by selecting **System Configuration** and then **Hypervisor Configuration**. The hypervisor mode shows a setting of either PowerVM or OPAL.

If OPAL is the wanted configuration, then you must remove this connection from the HMC by selecting **Connections** and then **Reset or Remove Connections**. Next, select **Remove Connections** and click **OK**.

Note: The OPAL hypervisor is not supported on the HMC.

If PowerVM is the wanted configuration, select **PowerVM** from the hypervisor mode menu and click **Continue**.

Note: The setting can be changed only when the server is powered off. To power off the server select **Power/Restart Control** and then **Power On/Off System**. Click **Save Settings and Power off** .

- **Connection not allowed**

  A connection state of **Version mismatch** and a reference code of **Connection not allowed 0009-0008-00000000** is returned when the FSP firmware and HMC versions are incompatible.

  To recover from this state, install an HMC version that supports the managed server model.

For more information about correction a **Version mismatch** state, see Version mismatch errors.

Systems Management for Servers

Systems Management displays tasks to manage servers, logical partitions, and frames. Use these tasks to set up, configure, view current status, troubleshoot, and apply solutions for servers.

These tasks are listed when a managed system is selected. The tasks listed in the menu pod change as selections are made in the work area.

System content pane

View and monitor the state, health, and capacity information of all the systems that are connected to the management console.

The content pane, in the middle portion of the window, displays all the available systems and the associated information for each server. You can choose to display the information in a table view or a gallery view.

Each system displays the current state of the system, the number of central processing units (CPUs) that are in use, CPUs that are available, the amount of random access memory (RAM) that is in use, and the RAM that is available. Additionally, if you choose a tabular format view, you can filter the information to be displayed by selecting the drop-down arrow in the upper-right corner of the table. Select the check box corresponding to the field that you want to display on the table. If you are using HMC Version 9.1.930, in the **All Systems** table, you can also view information about the activated and deferred firmware levels.

You can click the **properties** icon to display the following information:

- Current state
- Reference code
- Machine type
- Serial number
- System location
- Firmware level
- Group tags
- Attention LED

You can click the **capacity** icon to display the following information:

- Date of collection.
- Processor usage (installed, activated, available, average, and peak). The bar graph and the numerical value show the average usage (average divided by activated) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by activated). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Memory allocation (installed, activated, available, average, and peak). The bar graph and the numerical value show the average usage (average divided by activated) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by activated). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Network I/O usage (sent and received in kilobytes per second and in packets per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Storage I/O usage (written and read information in kilobytes per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Data collection.

You can hover over the systems in the **All systems** window to view the system model description.

Operations

Operations contains the tasks for operating managed systems.

Power Off

Shut down the managed system. Powering off the managed system will make all partitions unavailable until the system is again powered on.

Before you power off the managed system, ensure that all logical partitions have been shut down and that their states have changed from `Running` to `Not Activated`. For more information on shutting down a logical partition, see “Shut Down” on page 59

If you do not shut down all logical partitions on the managed system before you power off the managed system, the managed system shuts down each logical partition before the managed system itself powers off. This can cause a substantial delay in powering off the managed system, particularly if the logical partitions are not responsive. Further, the logical partitions might shut down abnormally, which could result in data loss and further delays when you activate the logical partitions once more.

Choose from the following options:

Normal power off
The Normal power off mode shuts down the system's operations in a controlled manner. During the shutdown, programs running active jobs are allowed to perform cleanup (end-of-job processing).

Fast power off
The Fast power off mode shuts down the system by stopping all active jobs immediately. The programs running those jobs are not allowed to perform any cleanup. Use this option when you need to shut down the system because of an urgent or critical situation.

Power On

Use the **Power On** task to start a managed system.

Choose from the following options to power on your managed system:

Normal: Select this option to specify that the HMC uses the current setting for the partition start policy to determine how to power on the managed system. The current setting can be one of the following values:

- **Auto-Start Always**: This option specifies that the HMC power on logical partitions automatically after the managed system powers on. If powering on the managed system is the result of a user action, the HMC starts all partitions that are configured for automatic startup. If powering on the managed system is the result of an automatic recovery process, the HMC starts only those logical partitions that were running at the time the system is powered off. This option is always available for selection.
- **Stop at Partition Standby**: This option specifies that logical partition startup is in standby mode after the managed system powers on and the HMC does not start any logical partitions when the managed system powers on. If powering on the managed system is the result of an automatic recovery process and the HMC is used to start a logical partition, the HMC starts all logical partitions that were running at the time the system is powered off. This option is available for selection only when the firmware for the managed system does not support advanced IPL capabilities.
- **Auto-Start for Auto-Recovery**: This option specifies that the HMC power on logical partitions automatically only after the managed system powers on as the result of an automatic recovery process. This option is available for selection only when the firmware for the managed system supports this advanced IPL capability.
- **User-Initiated**: This option specifies that the HMC does not start any logical partitions when the managed system powers on. You must start logical partitions manually on the managed system by using the HMC. This option is available for selection only when the firmware for the managed system supports this advanced IPL capability.

You can set the partition start policy from the **Power On Parameters** page of the **Properties** task for the managed system.

System profile: Selecting this power-on option specifies that the HMC power on the system and its logical partitions based on a predefined system profile. When you select this power-on option, you must select the partition profile that you want the HMC to use to activate logical partitions on the managed system.

Hardware Discovery: Selecting this power-on option specifies that the HMC run the hardware discovery process when the managed system powers on. The hardware discovery process captures information about all I/O devices, in particular those devices that are not currently assigned to partitions. When you select the hardware discovery **power on** option for a managed system, the managed system is powered on into a special mode that performs the hardware discovery. After the Hardware Discovery process is complete, the system will be in Operating state with any partitions in the power-off state. The hardware discovery process records the hardware inventory in a cache on the managed system. The collected information is then available for use when you display data for I/O devices or when you create a system plan based on the managed system. This option is available only if the system is capable of using the hardware discovery process to capture I/O hardware inventory for the managed system.

Power Management

You can reduce the processor power consumption of the managed system by enabling power saver mode.

About this task

To enable power saver mode, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon, and then select **All Systems**.
2. Select the server for which you want to enable the power saver mode and click **Actions** > **View All Actions**.
3. Select **Power Management** under **Operations**.
4. Choose from any of the following Power Saver mode options:
   - **Disable All modes**: Disables the Power Saver mode. The processor clock frequency is set to a nominal value and the power that is used by the system remains at a nominal level.

- **Enable Static mode**: Reduces the power consumption by reducing the processor clock frequency and the voltage to fixed values. This option also reduces the power consumption of the system while still delivering predictable performance.
- **Enable Dynamic Performance mode**: Causes the processor frequency to vary based on processor use. During periods of moderate or high use, the processor frequency is set to the maximum value allowed, which might be higher than the nominal frequency. Additionally, the frequency is set to a value that is less than the nominal frequency during periods of low processor use.
- **Enable Maximum Performance mode**: Causes the processor frequency to be set at a fixed value that you can specify. You can set the maximum limit of the processor frequency and power consumption of the system.

Note: Enabling any of the power saver modes causes changes in the processor frequencies, changes in processor use, changes in power consumption, and varying performance.

Schedule Operations

Create a schedule for certain operations to be performed on the managed system without operator assistance.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule power on or off operations for a managed system.

The Scheduled Operations task displays the following information for each operation:

- The processor that is the object of the operation.
- The scheduled date
- The scheduled time
- The operation
- The number of remaining repetitions

From the **Scheduled Operations** window, you can perform the following tasks:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

You can schedule an operation to occur once or you can schedule it to repeat. You must provide the time and date that you want the operation to occur. If you want the operation to repeat, you are asked to select the following options:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)
- The total number of repetitions. (required)

The operations that you can schedule for the managed system include the following operations:

Activate on a System Profile
Schedules an operation on a selected system for scheduling activation of a selected system profile.

Backup Profile Data
Schedules an operation to back up profile data for a managed system.

Power Off Managed System
Schedules an operation for a system power off at regular intervals for a managed system.

Power On Managed System
Schedules an operation for a system power-on at regular intervals for a managed system.

Manage Utility CoD processors
Schedules an operation for managing how your Utility CoD processors are used.

Manage Utility CoD processor minute usage limit
Creates a limit for Utility CoD processor usage.

Modify a Shared Processor Pool
Schedules an operation for modifying a shared processor pool.

Move a partition to a different pool
Schedules an operation for moving a partition to a different processor pool.

Change power saver mode on a managed system
Schedules an operation for changing a managed system's power saver mode.

Monitor/Perform Dynamic Platform Optimize
Schedules an operation for performing dynamic platform optimization and for sending an email notification alert to a user.

To schedule operations on the managed system, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select one or more managed systems and click **Actions** > **Schedule Operations**.
3. From the **Scheduled Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, click **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range...**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
4. To close the window, click **Options** and then click **Exit**.

Launch ASM Interface

The Hardware Management Console (HMC) can connect directly to the Advanced System Management Interface (ASMI) for a selected system.

The ASMI is an interface to the service processor that allows you to manage the operation of the server, such as auto power restart, and to view information about the server, such as the error log and vital product data.

To connect to the Advanced System Management Interface, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. In the content area, select one or more managed systems and click **Actions** > **View All Actions** > **Launch Advanced System Management (ASM)**.

Rebuild

You can extract the configuration information from the managed system and rebuild the information on the Hardware Management Console (HMC).

This task does not disrupt the operation of the running server.

Rebuilding the managed system updates the information on the HMC about the managed system. Rebuilding the managed system is useful when the state of the managed system is `Incomplete`. The Incomplete state means that the HMC cannot gather complete information from the managed system about logical partitions, profiles, or resources.

Rebuilding the managed system is different from refreshing the **HMC** window. When the managed system is rebuilt, the HMC extracts the information from the managed system. You cannot start other tasks while the HMC rebuilds the managed system. This process can take several minutes.

Change Password

Change the Hardware Management Console (HMC) access password on the selected managed system.

After the password is changed, you must update the HMC access password for all other HMCs from which you want to access this managed system.

Enter the current password and then, enter a new password and verify it by entering it again.

Attention LED

View system attention LED information, light specific LEDs to identify a system component, and test all LEDs on a managed system.

The system provides several LEDs that help identify various components, such as enclosures or field replaceable units (FRUs), in the system. For this reason, they are called **Identify** LEDs. Individual LEDs are on or near the components. The LEDs are located either on the component itself or on the carrier of the component (for example, memory card, fan, memory module, or processor). LEDs are either green or amber. Green LEDs indicate either of the following states:

- Electrical power is present.
- Activity is occurring on a link. (The system might be sending or receiving information).

Amber LEDs indicate a fault or identify condition. If your system or one of the components on your system has an amber LED turned on or flashing, identify the problem and take the appropriate action to restore the system back to normal.

You can activate or deactivate the following types of identify LEDs:

Identify LED for an enclosure

If you want to add an adapter to a specific drawer (enclosure), you need to know the machine type, model, and serial number (MTMS) of the drawer. To determine whether you have the correct MTMS for the drawer that needs the new adapter, you can activate the LED for a drawer and verify that the MTMS corresponds to the drawer that requires the new adapter.

Identify LED for a FRU associated with a specified enclosure

If you want to attach a cable to a specific I/O adapter, you can activate the LED for the adapter that is a field replaceable unit (FRU), and then physically verify where to attach the cable. This step can be especially useful when you have several adapters with open ports.

You can deactivate a system attention LED or a logical partition LED. For example, you might determine that a problem is not a high priority and decide to repair the problem later. However, you want to be alerted if another problem occurs, so you must deactivate the system attention LED so that it can be activated again if another problem occurs.

Choose from the following options:

Turn Attention LED Off

From this task, you can deactivate the system attention LED.

Identify Attention LED

Displays the current Identify LED states for all the location codes that are contained in the selected enclosure. From this task, you can select a single location code or multiple location codes to operate against and activate or deactivate one or more LEDs by selecting the corresponding button.

Test Attention LED

Initiates an LED Lamp Test against the selected system. All LEDs activate for several minutes.

Connections

You can view the Hardware Management Console (HMC) connection status to service processors or frames, reset those connections, connect another HMC to the selected managed system, or disconnect another HMC.

If you select a managed system in the work area, the following tasks pertain to that managed system. If you select a frame, the tasks pertain to that frame.

Service Processor Status

View information about the status of the Hardware Management Console (HMC) connection to the service processors on the managed system.

About this task

To show the service processor connection status to the service processors on the managed system, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to view the service processor connection status and click **Actions** > **View All Actions** > **Service Processor Status**.

Reset or Remove Connections

Reset or remove a managed system from the Hardware Management Console (HMC) interface.

About this task

To reset or remove connections, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server that you want to reset or remove and click **Actions** > **Reset or Remove System Connection**.
3. Select one of the options from **Reset Connection** or **Remove Connection** and click **OK**.

Disconnect Another HMC

You can disconnect a connection between a selected Hardware Management Console (HMC) and the managed server.

About this task

To disconnect another HMC, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to disconnect another Management Console and click **Actions** > **View All Actions** > **Disconnect Another HMC**.
3. Select an HMC from the list and click **OK**.

System Templates

System templates contain configuration details for resources such as system properties, shared processor pools, reserved storage pool, shared memory pool, Host Ethernet adapters, and SR-IOV adapters. Many of the system settings that you previously configured by using separate tasks are available in the **Deploy System from Template** wizard. For example, you can configure the Virtual I/O Servers, virtual network bridges, and virtual storage settings when you use the wizard to deploy a system from a system template.

The template library includes predefined system templates, which contain configuration settings based on common usage scenarios. Predefined system templates are available for your immediate use.

You can also create custom system templates that contain configuration settings that are specific to your environment. You can create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration.

Deploy System from Template

You can deploy systems by using system templates that are available in the template library of the Hardware Management Console (HMC). The Deploy System from Template wizard guides you to provide target system-specific information that is required to complete the deployment of the selected system.

Create Partition from Template

You can create a partition by using partition templates that are available in the template library of the Hardware Management Console (HMC). The Create a Partition from Template wizard guides you through the deployment process and configuration steps.

Capture Configuration as Template

You can capture the configuration details of a running server and save the information as a custom system template by using the Hardware Management Console (HMC). This function is useful if you want to deploy multiple servers with the same configuration. If you want to use a predefined template, you do not need to complete this task.

To capture configuration as a template, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **View All Actions**
3. Click **Capture Configuration as Template with Physical I/O** or **Capture Configuration as Template without Physical I/O**.
4. Enter a template name and description, and then click **OK**.

Use the online Help if you need additional information about capturing the configuration as a template.

Updates

Display tasks to view system information, manage updates on your Hardware Management Console (HMC), or check system readiness.

View System Information

Display information on a selected system from the Hardware Management Console (HMC).

To view the network topology, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates** > **View System Information**.
3. Select a LIC repository from the list and click **OK**.
4. When you have completed this task, click **Close**.

Use the online Help if you need additional information for viewing system information of the HMC.

Change Licensed Internal Code

Change the Licensed Internal Code of a managed system by using your Hardware Management Console (HMC).

You can change the Licensed Internal Code for the current release or to a new release.

To change the Licensed Internal Code, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates**.
3. Select **Change Licensed Internal Code** > **for the Current Release** or **Change Licensed Internal Code** > **to a New Release**.

   Note: Click **Start Change Licensed Internal Code** wizard to start a guided update of managed system, power, and I/O Licensed Internal Code (LIC). Click **View System Information** to examine current LIC levels, including retrievable levels. Click **Select Advanced Features** to update the managed system and power the LIC with more options and more targeting choices.
4. Select an action from the list and click **OK**.
5. When you complete this task, click **Close**.

Use the online Help if you need additional information for changing the Licensed Internal Code of the HMC.

Check System Readiness

Check the readiness of the Licensed Internal Code of a selected system from the Hardware Management Console (HMC).

To check system readiness, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates** > **Check System Readiness**.
3. When you have completed this task, click **OK**.

Use the online Help if you need additional information for checking system readiness of the HMC.

SR-IOV Firmware Update

Update the driver firmware for SR-IOV adapters on your Hardware Management Console (HMC).

Note: The adapter must be in shared mode.

To update the firmware for SR-IOV adapters, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates** > **SR-IOV Firmware Update**.
3. Select and right-click an adapter or adapters to get the context menu.
4. Select the type of firmware update to start.

   Note: Either the adapter driver firmware can be updated or both the adapter driver and adapter firmware can be updated. During the update operation of the adapter or adapter driver firmware, configured logical ports on the adapter might experience a temporary disruption of network traffic. Each adapter can take between 2 - 5 minutes to update. Updates are performed serially.
5. When you have completed this task, click **Close**.

Use the online Help if you need additional information for updating the driver or firmware for SR-IOV adapters.

Legacy

You can view **legacy** tasks that are available on the Hardware Management Console (HMC).

If you select a managed system in the work area, the following **legacy** tasks pertain to that managed system.

Partition Availability Priority

Use this task to specify the partition-availability priority of each logical partition on this managed system.

The managed system uses partition-availability priorities when a processor fails. If a processor fails on a logical partition and unassigned processors are not available on the managed system, then the logical partition can acquire a replacement processor from logical partitions with a lower partition-availability priority. This task allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

You can change the partition availability priority for a partition by selecting a partition and by choosing an availability priority from the list.

Use the online Help if you need additional information about prioritizing partitions.

View Workload Management Groups

Display a detailed view of the workload management groups that you specify for the managed system.

Each group displays the total number of processors, processing units for partitions that use shared mode processing, and the total amount of memory that is allocated to the partitions in the group.

Manage System Profiles

A system profile is an ordered list of partition profiles that is used by the Hardware Management Console (HMC) to start the logical partitions on a managed system in a specific configuration.

When you activate the system profile, the managed system attempts to activate each partition profile in the system profile in the order specified. A system profile helps you to activate or change the managed system from one set of logical partition configurations to another.

You can create a system profile that has a partition profile with overcommitted resources. You can use the HMC to validate the system profile against the currently available system resources and against the total system resources. Validating your system profile ensures that your I/O devices and processing resources are not overcommitted, and it increases the likelihood that the system profile can be activated. The validation process estimates the amount of memory that is needed to activate all of the partition profiles in the system profile. A system profile can pass validation and yet not have enough memory to be activated.

Use this task to complete the following tasks:

- Create new system profiles.
- Create a copy of a system profile.
- Validate the resources that are specified in the system profile against the resources available on the managed system. The validation process indicates whether any of the logical partitions in the system profile are already active and whether the uncommitted resources on the managed system can meet the minimum resources that are specified in the partition profile.
- View the properties of a system profile. From this task, you can view or change an existing system profile.
- Delete a system profile.
- Activate a system profile. When you activate a system profile, the managed system attempts to activate the partition profiles in the order that is specified in the system profile.

Use the online Help if you need additional information about managing system profiles.

Manage Partition Data

A partition profile is a record on the HMC that specifies a possible configuration for a logical partition. When you activate a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

A partition profile specifies the wanted system resources for the logical partition and the minimum and maximum amounts of system resources that the logical partition can have. The system resources that are specified within a partition profile includes processors, memory, and I/O resources. The partition profile can also specify certain operating settings for the logical partition. For example, you can set a partition profile such that, when the partition profile is activated, the logical partition is set to start automatically the next time that you power on the managed system.

Each logical partition on a managed system that is managed by an HMC has at least one partition profile. You can create more partition profiles with different resource specifications for your logical partition. If you create multiple partition profiles, you can designate any partition profile on the logical partition to be the default partition profile. The HMC activates the default profile if you do not select a specific partition profile to be activated. Only one partition profile can be active at one time. To activate another partition profile for a logical partition, you must shut down the logical partition before you activate the other partition profile.

A partition profile is identified by partition ID and profile name. Partition IDs are whole numbers that are used to identify each logical partition that you create on a managed system, and profile names identify the partition profiles that you create for each logical partition. Each partition profile on a logical partition must have a unique profile name, but you can use a profile name for different logical partitions on a single managed system. For example, logical partition 1 cannot have more than one partition profile with a profile name of normal, but you can create a profile named normal for each logical partition on the managed system.

When you create a partition profile, the HMC shows you all of the resources available on your system. The HMC does not verify whether another partition profile is using a portion of these resources. Therefore, it is possible for you to overcommit resources. When you activate a profile, the system attempts to allocate the resources that you assigned to the profile. If you overcommit resources, the partition profile is not activated.

For example, you have four processors on your managed system. Partition 1 profile A has three processors, and partition 2 profile B has two processors. If you attempt to activate both of these partition profiles at the same time, partition 2 profile B fails to activate because you overcommitted processor resources.

When you shut down a logical partition and reactivate the logical partition by using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition by using dynamic logical partitioning are lost when you reactivate the logical partition that uses a partition profile. This action is required when you want to undo dynamic logical partitioning changes for the logical partition. However, this action is not required if you want to reactivate the logical partition that uses the resource specifications that the logical partition had when you shut down the managed system. Therefore, keep your partition profiles up to date with the latest resource specifications. You can save the current configuration of the logical partition as a partition profile. This task avoids having to change partition profiles manually.

If you shut down a logical partition whose partition profiles are not up to date, and if the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system by using the partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

Use the Manage Partition Data tasks to complete the following tasks:

- Restore partition data. If you lose partition profile data, use the restore task in one of the following ways:
  - Restore partition data from a backup file. Profile modifications that are completed after the selected backup file was created are lost.
  - Restore merged data from your backup file and recent profile activity. The data in the backup file takes priority over recent profile activity if the information conflicts.
  - Restore merged data from recent profile activity and your backup file. The data from recent profile activity takes priority over your backup file if the information conflicts.
- Initialize partition data. Initializing the partition data for a managed system deletes all of the currently defined system profiles, partitions, and partition profiles.
- Back up a partition profile to a file.
- Back up partition data to a file.

Use the online Help if you need additional information about managing partition data.

Utilization Data

You can set the Hardware Management Console (HMC) to collect resource utilization data for a specific managed system or for all systems the HMC manages.

The HMC collects utilization data for memory and processor resources. You can use this data to analyze trends and make resource adjustments. The data is collected into records that are called events. Events are created at the following times:

- At periodic intervals (30 seconds, 1 minute, 5 minutes, 30 minutes, hourly, daily, and monthly).
- When you make system-level and partition-level state and configuration changes that affect resource utilization.

- When you start, shut down, and change the local time on the HMC.

You must set the HMC to collect utilization data for a managed system before utilization data can display for the managed system.

Use the **Change Sampling Rate** task to enable, set and change the sampling rate, or to disable sampling collection.

Create Partition

You can quickly create partitions with minimum resources.

To create a partition, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to create a partition and click **Actions** > **View System Partitions**.
3. Click **Create Partition**.
4. Complete the required information in the **Basic Partition Configuration**, **Processor Configuration**, and **Memory Configuration** tabs. If you want to assign all the system resources to the partition, select the **Assign all system resources** check box.
5. To create multiple partitions, move the slider to the right and select the **Multiple Partitions View**.
6. To add a new partition definition, click the **(+)** sign located on the top of the partition table.
7. Select the added partition and complete the required information in the **Basic Partition Configuration**, **Processor Configuration**, and **Memory Configuration** tabs. In the **Basic Partition Configuration** tab, you can provide details about the number of partition instances that you want to create. You can create a maximum of 20 partition instances.
8. To remove an existing partition, select the partition that you want to remove and click the **(-)** sign.
9. Click **OK**.

Use the online Help if you need additional information about this task.

Note: If the managed system supports virtual serial number and if the managed system is not in an Enterprise Pool 2.0, the **Virtual Serial Number** can be specified in the **Basic Partition Configuration** tab.

When the firmware level is at FW950 and the managed system has logical partitions that are assigned with virtual serial numbers, the managed system cannot be added to an Enterprise Pool 2.0. Also, if the managed system is in an Enterprise Pool 2.0, virtual serial number cannot be assigned to the logical partitions.

Properties

Displays the properties of the selected managed system. This information is useful in system and partition planning and resource allocation.

To open the properties tasks that are available for your system, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Properties** and then select the properties task that you want to perform from the list.

General Settings

View or change the general and advanced settings for the managed system.

These properties include the following tabs:

General Properties
The **General Properties** tab displays the system's name, serial number, model and type, state, attention led state, service processor version, maximum number of partitions, assigned service partition (if designated), and power off policy information.

Migration
View the partition mobility properties and change the migration policy for inactive partitions on the managed system.

Power-On Parameters
From the **Power-On Parameters** tab, you can change the power-on parameters for the next restart by changing the values in the **Next Value** fields. These changes are only valid for the next managed system restart.

Advanced
The **Advanced** tab displays huge page memory capabilities on the managed system, including available huge page memory, configurable huge page memory, current page size, and current maximum huge page memory. To change memory allocation on systems with huge page table support, set the Requested huge page memory (in pages) field to the wanted memory. To change the requested value for huge page memory, the system must be powered off.

The **Barrier Synchronization Register (BSR)** option displays array information.

The **Processor Performance** option displays the TurboCore mode and the System Partition Processor Limit (SPPL). You can set the next TurboCore mode and the next SPPL value. The SPPL applies to both dedicated processor partitions and shared processor partitions.

The **Memory Mirroring** option displays the current mirroring mode and the current system firmware mirroring status. You can set the next mirroring mode. You can also start the memory optimization tool.

You can view the VTPM settings.

Processor, Memory, I/O

View or change the memory, processor, and physical I/O resource settings for the managed system.

These properties include the following tabs:

Processor
The **Processor** tab displays information about the processors of the managed system, which includes:

- installed processing units
- unconfigured processing units
- available processor units
- available with stealable processor units
- configurable processing units
- minimum number of processing units per virtual processor
- maximum number of shared processor pools

The **Available with stealable** field displays the information about the available processing units, which is the sum of the available processing units in the managed system and the number of stealable processing units.

The stealable processor units value is the sum of the processor resources that are assigned to all the powered off or hibernated partitions on the managed system.

Notes:

- The information about stealable processor units is available only when the managed system is in the standby state or in the operating state.
- If the managed system is licensed with Power® IFL processor and if the firmware level is at FW910, or later, the **Available (with stealable)** field is displayed.
- When a POWER9 system is licensed with some IFL processors, the tab also displays the information about the remaining processors that are available for running the AIX or IBM i partitions.

Memory

The **Memory** tab displays information about the memory of the managed system, which includes:

- installed memory
- unconfigured memory
- available memory
- available with stealable memory
- configurable memory
- memory region size
- current memory available for partition usage
- system firmware current memory

The **Available with stealable** field displays the information about the available memory, which is the sum of available memory in the managed system and the amount of stealable memory resources. The tab also displays the maximum number of memory pools that are available.

Note: The information about stealable memory resources is available only when the managed system is in the standby state or in the operating state.

Physical I/O adapters

The **Physical I/O Adapters** tab displays the physical I/O resources for the managed system. The assignment of I/O slots and partition, the adaptor-type, and the slot LP limit information are displayed. The physical I/O resources information is grouped by units.

- The **Adapter Description** column displays the physical description of each resource.
- The **Physical Location Code** column displays the physical location code of each resource.
- The **Owner** column displays who currently own the physical I/O. The value of this column can be any of the following values:
  - When a single root I/O virtualization (SR-IOV) adapter is in the shared mode, **Hypervisor** is displayed in this column.
  - When an SR-IOV adapter is in the dedicated mode, **Unassigned** is displayed when the adapter is not assigned to any partition as a dedicated physical I/O.
  - When an SR-IOV adapter is in the dedicated mode, the logical partition name is displayed when the adapter is assigned to any logical partition as a dedicated physical I/O.
- The **Bus Number** column displays the bus number of the resource.
- The **I/O Pools** button displays all of the I/O pools found in the system and the partitions that are participating in the pools.

PowerVM

You can use the PowerVM function on the Hardware Management Console (HMC) to manage the system-level virtualization capabilities of your IBM Power Systems servers.

You can use the PowerVM task to manage virtual resources that are associated with a system, such as configuring a Virtual I/O Server (VIOS), virtual networks, and virtual storage. You can manage the PowerVM functions at the managed system level in response to changes in workloads or to enhance performance.

The PowerVM functions include the following tasks:

- Managing Virtual I/O Servers
- Managing virtual networks
- Managing virtual storage
- Managing hardware virtualized I/O (SR-IOV adapters, host Ethernet adapters (HEAs), and host channel adapters (HCAs))
- Managing a reserved processor pool
- Managing shared processor pools
- Managing a shared memory pool

Use the online Help if you need additional information about managing PowerVM.

Capacity on Demand

Activate disabled processors or memory that is installed on your managed server.

You can use Capacity on Demand (CoD) to nondisruptively activate (no startup needed) processors and memory. Capacity on Demand also gives you the option to temporarily activate capacity to meet intermittent performance needs, to activate extra capacity on a trial basis, and to access capacity to support operations in times of need.

Capacity on Demand Functions

Learn about the different Capacity on Demand functions that are available for your system.

You can use Capacity on Demand (CoD) to nondisruptively activate (no startup needed) processors and memory. Capacity on Demand also gives you the option to temporarily activate capacity to meet intermittent performance needs, to activate extra capacity on a trial basis, and to access capacity to support operations in times of need.

The **Capacity on Demand Processor** functions include the following tasks:

- View processor settings
- CUoD (permanent) processor
  - View CUoD code information
- On/Off processor
  - Manage
  - View billing information
  - View capacity settings
  - View code information
- Utility processor
  - Manage
  - View capacity settings
  - View code information
  - View shared processor utilization
- Trial processor
  - Stop trial
  - View capacity settings
  - View code information

The **Capacity on Demand Memory** functions include the following tasks:

- View memory settings
- CUoD (permanent) memory

  - View CUoD code information
- On/Off memory
  - Manage
  - View billing information
  - View capacity settings
  - View code information
- Trial memory
  - Stop trial
  - View capacity settings
  - View code information

Use the online Help if you need additional information about Capacity on Demand functions.

Licensed Capabilities

View and edit the runtime capabilities that are supported by the managed system.

You can view which licensed capabilities are active on your managed system. To activate a new licensed capability, click **Enter Activation Code** and enter the activation code.

The licensed functions that are available on the managed system include the following capabilities:

- Active Memory Sharing Capable
- Live Partition Mobility Capable
- Micro-Partitioning® Capable
- PowerVM Partition Simplified Remote Restart Capable
- SR-IOV Capable (Logical Port Limit)
- Virtual I/O Server Capable
- Active Memory Expansion Capable
- Active Memory Mirroring for Hypervisor Capable
- Coherent Accelerator Processor Interface (CAPI)
- AIX Enablement for 256-Core Partition Capable
- Dynamic Platform Optimization Capable
- IBM i 5250 Application Capable

Use the online Help if you need additional information about licensed capabilities.

Serviceability

Problem Analysis on the HMC automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. Use the **Serviceable Events Manager** task to view specific events for selected systems. However, if you notice a problem occurred or you suspect a problem is affecting the system but Problem Analysis has not reported it to you, use the **Create Serviceable Event** task to report the problem to your service provider.

To open the serviceability tasks that are available for your system, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.

3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the serviceability task that you want to perform from the list.

Tasks Log

View all the tasks that are currently running or completed on the Hardware Management Console (HMC).

To view the tasks log, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Tasks Log**.
2. You can view the following tabs in the tasks log:
   - **Task name**: Displays the name of task.
   - **Status**: Displays the current state of the task (running or completed).
   - **Resource**: Displays the name of the resource.
   - **Resource type**: Displays the type of resource.
   - **Initiator**: Displays the name of the user that initiated the task.
   - **Start time**: Displays the time that the task was initiated.
   - **Duration**: Displays the amount of time that the task took to complete.

Use the online Help for additional information about viewing the tasks log.

Serviceability

Problem Analysis on the Hardware Management Console (HMC) automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. Use the **Serviceable Events Manager** task to view specific events for selected systems. However, if you notice a problem occurred or you suspect a problem is affecting the system but Problem Analysis has not reported it to you, use the **Create Serviceable Event** task to report the problem to your service provider.

To open the serviceability tasks that are available for your system, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the serviceability task that you want to perform from the list.

Serviceable Events Manager

Problems on your managed system are reported to the Hardware Management Console (HMC) as serviceable events. You can view the problem, manage problem data, call home the event to your service provider, or repair the problem.

To set the criteria for the serviceable events you to view, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceable events.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Click **Serviceable Events Manager**.

5. Provide event criteria, error criteria, and FRU criteria. If you do not want the results to be filtered, select **ALL**. Click **Refresh** to refresh the list of serviceable events based on the criteria filter values.

The **Serviceable Events Overview** window displays all of the events that match your criteria. The information that is displayed in the compact table view includes the following information:

- Problem Number
- PMH Number
- Reference Code - Click the Reference code to display a description of the problem reported and actions that can be taken to fix the problem.
- Status of the problem
- Last reported time of the problem
- Failing MTMS of the problem
- Originating HMC

The full table view includes more detailed information, including reporting partition ID, primary data event timestamp, duplicate count, notification type, first reported time, reporting name, reporting MTMS, firmware fix, and serviceable event text.

Select a serviceable event and use the **Action** menu to:

- **View Details**: Field-replaceable units (FRUs) associated with this event and their descriptions.
- **View Files**: View the files associated with the selected serviceable event.
- **View Reference Code Description**: View the description of the reference code associated with the selected serviceable event. The option will not be available if additional description is not available.
- **Call Home**: Report the event to your service provider.
- **Repair**: Start a guided repair procedure, if available.
- **Close Event**: After the problem is solved, add comments and close the event.
- **Add PMH Comment**: Add a comment to a PMH for a selected serviceable event. If a PMH number does not exist for a given problem, the Add PMH Comment option is not available.

Use the online Help if you need additional information on managing serviceable events.

Create Serviceable Event

This task reports problems that occurred on your Hardware Management Console (HMC) to the service provider (for example, the mouse does not work) or lets you test problem reporting.

Submitting a problem is dependent upon whether you have customized this Hardware Management Console to use the Remote Support Facility (RSF) and if it is authorized to automatically call for service. If so, the problem information and service request is sent to the service provider automatically with a modem transmission.

To report a problem on your Hardware Management Console, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Create Serviceable Event**.
3. From the **Create Serviceable Event** window, select a problem type from the list displayed.
4. Enter a brief description of your problem in the **Problem Description** input field and then click **Request Service**.

To test problem reporting from the **Report a Problem** window:

1. Select **Test automatic problem reporting** and enter *This is just a test* in the **Problem Description** input field.

2. Click **Request Service**. The problems are reported to the service provider for the Hardware Management Console. Reporting a problem sends to the service provider the information you provide on **Report a Problem** window, and machine information that identifies the console.

Use the online Help if you need additional information for reporting a problem or testing if problem reporting works.

Manage Dumps

Manage system, service processor, and power subsystem dumps for systems that are managed by the Hardware Management Console (HMC).

system dump
A collection of data from server hardware and firmware, either after a system failure or a manual request. Perform a system dump only under the direction of your next level of support or your service provider.

service processor dump
A collection of data from a service processor either after a failure, external reset, or manual request.

power subsystem dump
A collection of data from Bulk Power Control service processor. This process is only applicable to certain models of managed systems.

Use the **Manage Dump** task to complete the following tasks:

- Initiate a system dump, a service processor dump, or a power subsystem dump.
- Modify the dump capability parameters for a dump type before you initiate a dump.
- Delete a dump.
- Copy a dump to media.
- Copy a dump to another system by using file transfer protocol (FTP).
- Call home a dump by using the Call Home feature to transmit the dump back to your service provider, for example IBM Remote Support, for further analysis.
- View the offload status of a dump as it progresses.

Use the online Help if you need additional information for managing dumps.

Collect VPD

Copy Vital Product Data (VPD) to removable media.

The managed system has VPD that is stored internally. The VPD consists of information such as how much memory is installed, and how many processors are installed. These records can provide valuable information that can be used by remote service and service representatives so that they can help you keep the firmware and software on your managed system up to date.

Note: To collect VPD, you must have at least one operational partition. For more information, see Logical Partitioning.

The information in the VPD file can be used to complete the following types of orders for your managed system:

- Install or remove a sales feature.
- Upgrade or rollback a model.
- Upgrade or rollback a feature.

Using this task, this information can be sent to removable media (diskette or memory key) for use by you or your service provider.

Use the online Help if you need additional information for collecting VPD.

Type, Model, Feature

Edit or display the model, type, machine type model serial (MTMS), or configuration ID of an enclosure.

The MTMS value or configuration ID for an expansion unit might need to be edited during a replacement procedure.

Use the online Help if you need additional information for editing MTMS.

Hardware Operations

Add, exchange, or remove hardware from the managed system. Display a list of installed FRUs or enclosures and their locations. Select a FRU or an enclosure and launch a step-by-step procedure to add, exchange, or remove the unit.

To open the hardware tasks that are available for your system, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage hardware tasks.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the hardware operations task that you want to perform from the list.

Prepare for Hot Repair or Upgrade

Provides a summary of required actions to be performed to isolate a particular hardware component as part of a service procedure.

From the **Component List** table, you can select the component to be repaired using the location code on the system to be repaired as directed by an Authorized Service Representative.

Power On/Off Unit

Use the **Power On/Off Unit** task to power on or off an I/O unit.

Only units or slots that reside in a power domain can be powered on or off. The corresponding power on/off buttons are disabled for location codes that are not controllable by the HMC.

Add FRU

Locate and add a Field Replaceable Unit (FRU).

To add a FRU to a POWER9 system, complete the following steps:

1. Select an enclosure type from the **Enclosure** menu.
2. Select a FRU type from the displayed list of FRU types for this enclosure, and click **Next**.
3. Select a FRU location, then click **Next** to start the Add FRU procedure for the selected location.
4. Follow the **Procedure** details listed in the window. The **Operation** list displays the correct operations for the procedure that you want to execute. The procedures direct when to perform each operation. Wait until advised to perform the operation before selecting the operation in the **Operation** list. Click **Next** to continue to the next service procedure.

   Note: Select the FRU location image on the header to display the location of the part to be serviced. To see a larger version of the FRU location image in a separate window, click **Display FRU Location**.
5. Click **Finish** to end the service when you have completed the last service procedure.

When the managed system is a POWER8® or earlier, to add a FRU, complete the following steps:

1. Select an enclosure type from the **Add FRU** menu.
2. Select a FRU type from the menu.
3. Click **Next**.
4. Select a location code from the displayed menu.
5. Click **Add**.

6. Click **Launch Procedure**.
7. When you complete the FRU installation process, click **Finish**.

Exchange FRU
Use the **Exchange FRU** task to exchange one field replaceable unit (FRU) with another FRU.

When the managed system is a POWER9 or later, to exchange a FRU, complete the following steps:

1. Select an installed enclosure type from the **Enclosure** menu.
2. Select a FRU type to be replaced, from the displayed list of FRU types for this enclosure and click **Next**.
3. Select a installed FRU location, then click **Next** to start the Exchange / Replace FRU procedure for the selected FRU.
4. Follow the **Procedure** details listed in the window. The **Operation** list displays the correct operations for the procedure that you want to execute. The procedures direct when to perform each operation. Wait until advised to perform the operation before selecting the operation in the **Operation** list. Click **Next** to continue to the next service procedure.

   Note: Select the FRU location image on the header to display the location of the part to be serviced. To see a larger version of the FRU location image in a separate window, click **Display FRU Location**.
5. Click **Finish** when you complete the exchange procedure.

When the managed system is a POWER8 or earlier, to exchange a FRU, complete the following steps:

1. Select an installed enclosure type from the **Exchange FRU** menu.
2. From the displayed list of FRU types for this enclosure, select a FRU type.
3. Click **Next** to display a list of locations for the FRU type.
4. Select a location code for a specific FRU.
5. Click **Add** to add the FRU location to **Pending Actions**.
6. Select **Launch Procedure** to begin replacing the FRUs that are listed in **Pending Actions**.
7. Click **Finish** when you complete the installation.

Remove FRU
Use the **Remove FRU** task to remove a FRU from your managed system.

When the managed system is a POWER9 or later, to remove a FRU, complete the following steps:

1. Select an enclosure from the menu to display a list of FRU types that are currently installed in the selected enclosure.
2. Select a FRU type from the displayed list of FRU types available for removal from the selected system and click **Next**.
3. Select a FRU location, then click **Next** to start the Remove FRU procedure for the selected FRU .
4. Follow the **Procedure** details listed in the window. The **Operation** list displays the correct operations for the procedure you want to execute. The procedures direct when to perform each operation. Wait until advised to perform the operation before selecting the operation in the **Operation** list. Click **Next** to continue to the next service procedure.

   Note: Select the FRU location image on the header to display the location of the part to be serviced. To see a larger version of the FRU location image in a separate window, click **Display FRU Location**.
5. Click **Finish** when you complete the removal procedure.

When the managed system is a POWER8 or earlier, to remove a FRU, complete the following steps:

1. Select an enclosure from the menu to display a list FRU types that are currently installed in the selected enclosure.
2. From the displayed list of FRU types for this enclosure, select a FRU type.
3. Click **Next** to display a list of locations for the FRU type.
4. Select a location code for a specific FRU.

5. Click **Add** to add the FRU location to **Pending Actions**.
6. Select **Launch Procedure** to begin removing the FRUs listed in **Pending Actions**.
7. Click **Finish** when you complete the removal procedure.

Add Enclosure
Learn how to locate and add an enclosure.

To add an enclosure, complete the following steps:

1. Select an enclosure type, then click **Add**.
2. Click **Launch Procedure**.
3. When you complete the enclosure installation process, click **Finish**.

Remove Enclosure
Use the **Remove Enclosure** task to remove an enclosure.

To remove an enclosure, complete the following steps:

1. Select an enclosure type, then click **Add** to add the selected enclosure type's location code to **Pending Actions**.
2. Click **Launch Procedure** to begin removing the enclosures that are identified in **Pending Actions** from the selected system.
3. Click **Finish** when you complete the enclosure removal process.

Open MES
View MES order numbers and their states, for any MES operations active or inactive for the Hardware Management Console (HMC).

Use **Add MES Order Number** to add a new order number to the list. To add an order number, complete the following steps:

1. Click **Add MES Order Number**.
2. Enter new MES order number.
3. Click **OK**.

Close MES
Close MES order numbers.

Use **Close MES** to close a MES. To close a MES, complete the following steps:

1. Select an open MES order number from the table.
2. Click **OK**.

Setup FSP Failover
Set up a secondary service processor if your managed system's primary service processor fails.

FSP Failover is designed to reduce customer outages due to service processor hardware failures. If a redundant service processor is supported for the current system configuration, select **Setup** to set up FSP Failover for the selected managed system.

To set up the FSP failover, complete the following steps:

1. In the content pane under **FSP failover**, click **Setup**.
2. Click **OK** to enable automatic failover for the selected system.

Initiate FSP Failover
Initiate a secondary service processor if your managed system's primary service processor fails.

FSP Failover is designed to reduce customer outages due to service processor hardware failures. Select **Initiate** to start the FSP Failover for the selected managed system.

To start the FSP failover, complete the following steps:

1. In the content pane under **FSP failover**, click **Initiate**.
2. Click **OK** to start the automatic failover for the selected system.

Reference Code Log

Reference codes provide general diagnostic, troubleshooting, and debugging information.

View reference codes that are generated for the selected managed system. Reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem.

To view the reference code history, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **Reference Code Log**.
4. Select a specific reference code to view the details.

Use the online Help if you need additional information about this task.

RIO Configuration

View the current hardware topology and the last valid hardware topology.

Displays the current hardware and last valid hardware topology. Any discrepancies between the current topology and the last valid topology are identified as errors.

To view the hardware topology, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **RIO Configuration**.
4. View the hardware topology information.

Use the online Help if you need additional information about this task.

PCI Configuration

View information about the Peripheral Component Interconnect Express (PCIe) hardware topology.

The PCIe hardware topology utility provides information about the PCIe links that exist for each system.

To view the PCIe hardware topology, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **PCI Configuration**.
4. View the PCIe hardware topology.

Use the online Help if you need additional information about this task.

Topology diagrams

Learn how to view the topology diagrams of a partition.

You can use the Hardware Management Console (HMC) to view the topology diagrams of a partition.

Viewing virtual networking diagrams

You can view the end-to-end network configuration for the selected system, by using the Hardware Management Console (HMC). The view of the virtual networks begins with the physical adapter cards and the physical ports that are connected to them. As you scroll down, you can see the defined virtual bridges, link aggregation devices, virtual switches, virtual networks, and partitions in the VIOS.

You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the network diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the end-to-end network configuration for the selected system by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Topology** and then click **Virtual Networking Diagram**.
4. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
5. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
6. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols that are used in the virtual networking diagram.

Use the online Help if you need additional information about this task.

Viewing virtual storage diagrams

Two types of virtual storage diagrams are available; systems storage and partition storage. You can view the virtual storage configuration for the selected system, including the physical and virtual components of system storage, by using the HMC. You can also view the virtual storage configuration for a single partition in a particular system, including the physical and virtual components of storage assigned to that particular partition, by using the Hardware Management Console (HMC).

This diagram displays a high-level overview of the contents of the system or a single partition, and not specific component relationships. You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the storage diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the virtual storage configuration for the selected system or a single partition by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Topology** and then click **Virtual Storage Diagram**.

   Note: To view the virtual storage diagrams of a single partition in a particular system, select the partition of your choice and then expand **Topology** and click **Partition Virtual Storage Diagram**
4. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
5. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
6. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols that are used in the virtual storage diagram.

Use the online Help if you need additional information about this task.

Viewing SR-IOV and vNIC diagrams

You can view the SR-IOV and virtual Network Interface Controllers (vNIC) configuration for the selected system, including the physical and virtual components, by using the Hardware Management Console (HMC).

This diagram displays the relationships between SR-IOV adapters and other virtual components, such as vNIC. You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the SR-IOV and vNIC diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the end-to-end network configuration for the selected system by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Topology** and then click **SR-IOV vNIC Diagram**.
4. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
5. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
6. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols that are used in the SR-IOV and vNIC diagram.

Use the online Help if you need additional information about this task.

Systems Management for Partitions

Systems Management displays tasks that you can perform to manage servers, logical partitions, and frames. Use these tasks to set up, configure, view current status, troubleshoot, and apply solutions for partitions.

The following sets of tasks are represented when a partition is selected and is shown in the menu pod or content pane. The tasks that are listed in the menu pod change as selections are made in the work area.

Partition content pane

View and monitor the state, health, and capacity information of all the partitions that are connected to the management console.

The content pane, in the middle portion of the window, displays all the available partition and the associated information for each partition.

Each partition displays the current state of the partition, the reference code, the number of virtual processors that are allocated, and the amount of random access memory (RAM) that is allocated. Additionally, if you choose a tabular format view, you can filter the information to be displayed by selecting the drop-down arrow in the upper-right corner of the table. Select the check box corresponding to the field that you want to display on the table. If you are using HMC Version 9.1.930, or later, in the **All partition** table, you can also view the memory mode of all partitions that are associated with the managed system.

You can click the **properties** icon to display the following information:

- Current state
- System name
- Reference code
- Partition ID
- IP address
- Environment
- OS version
- RMC connection
- Last activated profile
- Contains physical I/O
- Group tags
- Attention LED

You can click the **capacity** icon to display the following information:

- Date of collection.
- Processor usage (type (dedicated, uncapped, or capped), entitled capacity, and virtual processors). When the processor type is dedicated, the bar graph and the numerical value show the used processor usage (used divided by assigned) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by assigned). When the processor type is uncapped, the bar graph and the numerical value show the used processor usage (used divided by the number of virtual processors) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by the number of virtual processors). When the processor type is capped, then the bar graph and the numerical value show the used processor usage (used divided by entitled) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by entitled). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Memory allocation (allocated).
- Network I/O usage (sent and received in terabytes per second and in packets per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the

number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.

- Storage I/O usage (written and read information in kilobytes per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Data collection.

Partition Properties

The **View Partition Properties** task displays the selected partition's properties. This information is useful in resource allocation and partition management. These properties include:

General

The **General** tab displays the partition's name, ID, environment, state, resource configuration, operating system, boot mode to start the operating system, and the system on which the partition is located.

Note: If the managed system supports virtual serial number and if the managed system is not in an Enterprise Pool 2.0, the Virtual Serial Number property is displayed in the **General** tab.

Click **Advanced settings** to also view the list of supported hardware accelerators for a logical partition and Quality of Service (QoS) credits for a specific hardware accelerator. This section is not displayed if the managed system does not support hardware accelerators.

Note: When the HMC is at V9.2.950.0, or later, and when the firmware is at level FW950, or later, the **KeyStore Size** value can be chosen in the range 4 KB - 64 KB as the keystore size of the logical partition. The value of 0 KB indicates that the keystore function is disabled for the logical partition.

Processor

The **Processor** tab displays the current usage of processors.

Note: When the operating system and the hypervisor support a minimum entitlement of 0.05 processor per virtual processor, the minimum, maximum, and desired processing units can be set to the lowest supported value of 0.05.

Memory

The **Memory** tab displays properties of the running logical partition that is using the dedicated or the shared memory.

Physical I/O adapter

The **Physical I/O Adapter** tab displays the properties of all the physical I/O adapters that are available for the managed system and that can be assigned to a partition. You can also add and remove an adapter in a partition.

Change Default Profile

Change the default profile for the partition.

Select a profile from the drop down list to be the new default profile.

Operations

Operations contains the tasks for operating partitions.

About this task

To open the operations tasks that are available for your partitions, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Partitions**.
2. Select the partition for which you want to manage operations tasks. Click **Actions** > **View Partition Properties**
3. In the menu pod, expand **Partition Actions** and then expand **Operations**.
4. Select the operations task that you want to perform from the list.

Activate

Use the **Activate** task to activate a partition on your managed system that is in the **Not Activated** state.

Select the partition profile from the list of profiles and click **OK** to activate the partition. On the **Advanced** tab, select the **No VSI Profile** check box to ignore the failure while configuring the Virtual Station Interface (VSI) profile.

Note: As of HMC Version 7.7, or later, you can install a Virtual I/O Server (VIOS) on a logical partition from an HMC by using a DVD, a saved image, or a Network Installation Management (NIM) server.

Netboot

Use the **netboot** task to network boot an AIX, Linux, or an IBM i partition on your managed system that is in the **Not Activated** state.

The **Network boot** wizard guides you through the steps of installing the operating system on the partition and then activating the partition. Select a partition profile to install the operating system on the partition. Click **Next** to configure the network settings for the logical partition.

Note: For Virtual I/O Server, you must choose the **Install** option from the **Actions** menu to install the VIOS on your managed system that is in the **Not Activated** state.

Restart

Restart the selected logical partition or partitions.

For IBM i logical partitions, use this window only if you cannot restart the IBM i logical partition from the command line of the operating system. Using this window to restart an IBM i logical partition results in an abnormal IPL.

If you choose to restart VIOS partitions that are acting as the Paging Service Partition (PSP) for a number of client partitions, a warning displays, indicating that you must shut down the client partitions before you shut down the VIOS partition.

Choose one of the following options. The Operating System option and the Operating System Immediate option are enabled only if Resource Monitoring and Control (RMC) is up and configured.

Dump

The HMC shuts down the logical partition and initiates a main storage or system memory dump. For AIX and Linux logical partitions, the HMC also notifies the logical partition to shut down. For IBM i logical partitions, the processors are stopped immediately. After the shutdown is complete, the logical partition is immediately restarted. (IBM i logical partitions are restarted multiple times so that the logical partition can store the dump information.) Use this option if a portion of the operation system appears hung and you want a dump of the logical partition for analysis.

Operating System

The HMC shuts down the logical partition normally by issuing a shutdown -r command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. After the shutdown is complete, the logical partition is immediately restarted. This option is only available for AIX logical partitions. Immediate: The HMC shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs that are running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data is partially updated. Use this option only after a controlled end is unsuccessfully attempted.

Operating System Immediate

The HMC shuts down the logical partition immediately by issuing a shutdown -Fr command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. After the shutdown is complete, the logical partition is immediately restarted. This option is only available for AIX logical partitions.

Dump Retry

The HMC retries a main storage or system memory dump on the logical partition. After this operation is complete, the logical partition is shut down and restarted. Use this option only if you previously tried the **Dump** option without success. This option is only available for IBM i logical partitions.

Shut Down

Shut down the selected logical partition or partitions.

For IBM i logical partitions, use this window only if you cannot shut down the IBM i logical partition from the command line of the operating system. Using this window to shut down an IBM i logical partition will result in an abnormal IPL.

If you choose to shut down VIOS partitions that are acting as the Paging Service Partition (PSP) for a number of client partitions, a warning displays, indicating that you should shut down the client partitions before shutting down the VIOS partition.

Choose from the following options:

Delayed

The HMC shuts down the logical partition using the delayed power-off sequence. This allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it will end abnormally and the next restart may be longer than normal.

Immediate

The HMC shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

Operating System

The HMC shuts down the logical partition normally by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

Operating System Immediate

The HMC shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

Delete

Use the **Delete** task to delete the selected partition.

The Delete task deletes the selected partition and all of the partition profiles associated with the partition from the managed system. When you delete a partition, all hardware resources currently assigned to that partition become available to other partitions.

Schedule Operations

Create a schedule for certain operations to be performed on the logical partition.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule an operation to remove resources from a logical partition or move resources from one logical partition to another.

The Scheduled Operations task displays the following information for each operation:

- The processor that is the object of the operation
- The scheduled date
- The scheduled time
- The operation
- The number of remaining repetitions

From the **Scheduled Operations** window you can perform the following operations:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

You can schedule an operation to occur once or you can schedule it to repeat. You must provide the time and date that you want the operation to occur. If you want the operation to repeat, you are asked to select the following time intervals:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)
- The total number of repetitions. (required)

The operations that you can schedule for a logical partition include the following operations:

Activate on an LPAR
Schedules an operation on a selected profile for activation of the selected logical partition.

Dynamic Reconfiguration
Schedules an operation for adding, removing, or moving a resource (processors or megabytes of memory).

Operating System Shutdown (on a partition)
Schedules a shutdown of the selected logical partition.

To schedule operations on the HMC, complete the following steps:

1. In the Navigation area, click **Systems Management**.
2. In the work pane, select one or more partitions.
3. In the taskpad, select the **Operations** task category, then click **Schedule Operations**. The **Customize Scheduled Operations** window opens.
4. From the **Customize Scheduled Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, click **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.

5. To return to the HMC workplace, point to **Operations** and then click **Exit**.

Validate Maintenance Readiness

Use the **Validate Maintenance Readiness** task to validate the readiness of the Virtual I/O Server (VIOS) for maintenance. The VIOS must be in **Running** state with an active Resource Monitoring Control (RMC) connection to perform the validation operation on the VIOS. To complete the validation operation, you must have access to all the partitions of the managed system.

The Hardware Management Console (HMC) validates the readiness of the VIOS for the maintenance. When you execute the maintenance readiness operation, the HMC validates all the client logical partition that use Virtual I/O Servers for Multi-path I/O operation or redundancy setup for the network and storage that is attached to a logical partition. To check the redundancy setup of the network or storage, the HMC gets the inventory information of other Virtual I/O Servers that are associated with the managed system. However, if other VIOS partitions in the system do not have a proper RMC connection, the validation process continues, and results are shown based on the current states of the Virtual I/O Servers.

The page also displays information about all the impacted client partitions that do not have a redundant Virtual SCSI Storage, Virtual Fibre Channel, Virtual networks, and Virtual NIC that is provided by the VIOS.

Mobility

Use the Mobility task to migrate your partition to another server, ensure that the requirements for the migration are met, and recover if the partition is in an invalid state.

Migrate

Migrate a partition to another managed system.

About this task

To migrate a partition to another system, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. In the content pane, select the server. Click **Actions** > **View System Partitions**.
3. In the content pane, select the partition that you want to migrate to another system.
4. Click **Actions** > **Mobility** > **Migrate**. The Partition Migration wizard opens.
5. Complete the steps in the Partition Migration wizard and click **Finish**.

Validate

Validate the settings for moving the partition from the source system to the destination system.

About this task

To validate the settings, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. In the content pane, select the server. Click **Actions** > **View System Partitions**.
3. In the content pane, select the partition for which you want to validate the settings.
4. Click **Actions** > **Mobility** > **Validate**. The Partition Migration Validation window opens.
5. Fill in the information in the fields, and click **Validate**.

Recover

Recover this partition from a migration that did not complete.

About this task

To recover this partition from a migration that did not complete, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Systems**.
2. In the content pane, select the server. Click **Actions** > **View System Partitions**.
3. In the content pane, select the partition that you want to recover.
4. Click **Actions** > **Mobility** > **Recover**. The Migration Recovery window opens.
5. Complete the information as necessary and click **Recover**.

Partition Templates

Partition templates contain details for partition resources, such as physical adapters, virtual networks, and storage configuration. You can create client partitions from the quick-start templates that are available in the template library or from your own user-defined templates on the Hardware Management Console (HMC).

Capture Partition as a Template

You can capture the configuration details of a running partition and save the information as a partition template by using the Hardware Management Console (HMC).

To capture the configuration as a template, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Partitions**.
2. Select the partition for which you want to capture as a template.
3. Click **Actions** > **Templates** > **Capture Partition as a Template**.
4. Enter a template name and description.
5. Click **OK**.

Use the online Help if you need additional information about this task.

Profiles

Learn about the tasks that are available in the **Profiles** menu.

Manage Profiles

Use the **Manage Profiles** task to create, edit, copy, delete, or activate a profile for the selected partition.

A partition profile contains the resource configuration for the partition. You can modify the processor, memory, and adapter assignments for a profile by editing the profile.

The default partition profile for a logical partition is the partition profile that is used to activate the logical partition if no other partition profile is selected. You cannot delete the default partition profile unless you first designate another partition profile as the default partition profile. The default profile is defined in the status column.

Choose **Copy** to create an exact copy of the selected partition profile. This allows you to create multiple partition profiles that are nearly identical to one another by copying a partition profile and changing the copies as needed.

Manage Custom Groups

Groups are comprised of logical collections of objects. You can report status on a group basis, allowing you to monitor your system in a way that you prefer. You can also nest groups (a group contained within a group) to provide hierarchical or topology views.

One or more user-defined groups might already be defined on your Hardware Management Console (HMC). Default groups are listed under **Custom Groups** node under **Configuration**. The default groups are **All Partitions** and **All Objects**. You can create others, delete the ones that were created, add to created groups, create groups using the pattern match method, or delete from created groups by using the **Manage Custom Groups** task.

Use the online Help if you need additional information for managing custom groups.

Save Current Configuration

Save the current configuration of a logical partition to a new partition profile by entering a new profile name.

This procedure is useful if you change the configuration of a logical partition using dynamic logical partitioning and you do not want to lose the changes when you restart the logical partition. You can perform this procedure at any time after you initially activate a logical partition.

Delete Partition

You can delete a partition and the associated partition profile by using the Hardware Management Console (HMC).

To delete a partition, complete the following steps:

1. In the navigation area, click the **Resources** icon, and then select **All Partitions**.
2. Select the partition for which you want to delete.
3. Click **Actions** > **Delete Partition**.
4. Select any options that you want.
5. Click **OK**.

Use the online Help if you need additional information about this task.

Serviceability

Problem Analysis on the HMC automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. Use the **Serviceable Events Manager** task to view specific events for selected systems. However, if you notice a problem occurred or you suspect a problem is affecting the system but Problem Analysis has not reported it to you, use the **Create Serviceable Event** task to report the problem to your service provider.

Serviceable Events Manager

Problems on your managed partitions are reported to the Hardware Management Console (HMC) as serviceable events. You can view the problem, manage problem data, call home the event to your service provider, or repair the problem.

To set the criteria for the serviceable events you to view, complete the following steps:

1. In the navigation area, click the **Resources** icon, and then select **All Systems**.
2. Select the server for which you want to manage serviceable events.

3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Click **Serviceable Events Manager**.
5. Provide event criteria, error criteria, and FRU criteria. If you do not want the results that are filtered, select **ALL**. Click **Refresh** to refresh the list of serviceable events based on the criteria filter values.

The **Serviceable Events Overview** window displays all of the events that match your criteria. The information that is displayed in the compact table view includes the following information:

- Problem Number
- PMH Number
- Reference Code - Click the Reference code to display a description of the problem reported and actions that can be taken to fix the problem.
- Status of the problem
- Last reported time of the problem
- Failing MTMS of the problem
- Originating HMC

The full table view includes more detailed information, including reporting partition ID, primary data event timestamp, duplicate count, notification type, first reported time, reporting name, reporting MTMS, firmware fix, and serviceable event text.

Select a serviceable event and use the **Action** drop-down menu to:

- **View Details**: Field-replaceable units (FRUs) associated with this event and their descriptions.
- **View Files**: View the files associated with the selected serviceable event.
- **View Reference Code Description**: View the description of the reference code associated with the selected serviceable event. The option will not be available if additional description is not available.
- **Call Home**: Report the event to your service provider.
- **Repair**: Start a guided repair procedure, if available.
- **Close Event**: After the problem is solved, add comments and close the event.
- **Add PMH Comment**: Add a comment to a PMH for a selected serviceable event. If a PMH number does not exist for a given problem, the Add PMH Comment option is not available.

Use the online Help if you need additional information on managing serviceable events.

Reference Code Log

Use the **Reference Code Log** task to view reference codes that have been generated for the selected logical partition. Reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem.

By default, only the most recent reference codes that the logical partition has generated are displayed. To view more reference codes, enter the number of reference codes that you want to view into **View history** and click **Refresh**. The window displays that number of the latest reference codes, with the date and time at which each reference code was generated. The window can display up to the maximum number of reference codes stored for the logical partition.

Control Panel Functions

This task displays the available virtual control panel functions for the selected IBM i partition. The tasks are:

(21) Activate Dedicated Service Tools
Starts Dedicated Service Tools (DST) on the partition.

(65) Disable Remote Service
Deactivates remote service on the partition.

(66) Enable Remote Service
Activates remote service on the partition.

(68) Concurrent Maintenance Power Off Domain
Concurrent maintenance power domain Power Off.

(69) Concurrent Maintenance Power On Domain
Concurrent maintenance power domain Power On.

Virtual I/O

Learn how to view the virtual networks, virtual network interface controllers, and virtual storage of a partition.

You can use the Hardware Management Console (HMC) to view the virtual topology of a partition.

Virtual Networks

You can view and add virtual networks that are associated with the selected logical partition.

The **Virtual Networks** table lists the virtual network name, VLAN ID, virtual switch, virtual network bridge, and virtual Ethernet adapter ID that are associated with each virtual network. You can click **Attach Virtual Network** to view the available virtual networks and attach additional virtual networks to the logical partition.

To view the virtual networks for the selected partitions by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Virtual Networks**.

Use the online Help if you need additional information about this task.

Virtual NIC

You can manage all aspects of the virtual Network Interface Controller (NIC) configuration that is associated with the partition.

A virtual NIC is a type of virtual adapter that can be configured on logical partitions to provide a network interface. Each virtual NIC client adapter is backed by an SR-IOV logical port that is owned by the hosting partition.

The **Virtual NIC** table lists all virtual NICs that are configured for the selected partition. A virtual NIC can have one or more backing devices. The maximum number of backing devices per virtual NIC depends on the system. If the virtual NIC has more than one backing device, you can expand the node to view all the backing devices. If the virtual NIC has only one backing device, that backing device is the active backing device. The active backing device is the one that is in use by the virtual NIC. If the managed system is not failover capable, the table displays virtual NICs that have a single backing device.

You can add a virtual NIC to the partition. To add a virtual NIC, click **Add Virtual NIC**. You can configure the virtual NIC only in dedicated mode. You can also modify and view virtual NIC properties. To modify properties of a virtual NIC, select the virtual NIC in the table and click **Action** > **Modify vNIC** . To view the properties of a virtual NIC, select the virtual NIC in the table and click **Action** > **View vNIC**.

To view the virtual NIC for the selected partition by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.

3. In the menu pod, expand **Virtual I/O** and then click **Virtual NICs**.

Use the online Help if you need additional information about this task.

Virtual Storage

You can create, view, and manage the storage capability of the logical partition.

The **Virtual Storage** table displays the Virtual Small Computer Serial Interface (SCSI) devices that are configured to a logical partition. You can also view the information about the physical volume groups, shared storage pool volume, and the logical volume.

You can add the virtual storage resources to a partition. Click **Adapter View** to create, view the adapter configuration of the virtual storage devices that are allocated for the logical partition. Click **Storage View** to view and manage the storage capability of the logical partition.

Physical volumes can be exported to partitions as virtual SCSI disks. Click **Show assigned physical volumes** to view the assigned physical volumes that are assigned to the logical partition.

To add physical volumes to a partition, select the physical volumes from the list and specify the **User Defined Name** for each physical volume that you want to add to the partition and then click **OK**. If you want to change the server adapter ID that is assigned to each physical volume, click **Edit** for each of the physical volumes that you want to update. The **Edit connection** window is displayed. You can specify up to 3 Virtual I/O servers, and then enter the new server adapter ID that you want to assign for the adapter connection.

To add different types of virtual storage devices to a partition, click **Add Virtual SCSI Device**. Select the available virtual storage that you want to add. You can select the virtual storage types such as **Physical Volume**, **Shared Storage Pool Volume**, or **Logical Volume**.

To view the virtual storage for the selected partition by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon, and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Virtual Storage**.

Use the online Help if you need additional information about this task.

Hardware Virtualized I/O

You can view and change the settings of hardware virtualized I/O adapters, such as single root I/O virtualization (SR-IOV) port adapters and logical host Ethernet adapters (LHEA) for a partition by using the Hardware Management Console (HMC).

To view the hardware virtualized I/O adapters for the selected partition by using the HMC, complete the following steps:

1. In the navigation area, click the **Resources** icon, and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Hardware I/O**.
4. In the **SRIOV** tab, you can add an SR-IOV logical port to the partition or change the settings of the SR-IOV logical ports. In the **SR-IOV logical port** table, you can also view the information about the logical ports that can be migrated and the information about the backing device that is configured for the logical ports. In the **Logical host Ethernet adapters** (LHEA) tab, you can change the settings of an LHEA adapter. You can also add and remove an LHEA adapter.

Notes:

- With HMC Version 9.1.930, or later, the HMC also supports the RDMA over Converged Ethernet (RoCE) adapter.
- If you are using HMC Version 9.1.940, with firmware at level FW940, or later, you can create logical partitions that have an SR-IOV logical port that can be migrated. You can migrate a logical partition with SR-IOV logical ports when the Migratable option is used to create a backup virtual device when creating a logical port. The backup device can be either a virtual Ethernet or a virtual Network Interface Controller (NIC) adapter. When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Use the online Help if you need additional information about this task.

Systems Management for Frames

Set up, configure, view status, troubleshoot, and apply solutions for frames.

Properties

Display the selected frame properties.

Frame properties include the following properties:

General
The **General** tab displays the frame name and number, state, type, model, and serial number.

Managed Systems
The **Managed Systems** tab displays all of the managed systems that are contained in the frame and their cage numbers. A cage is a division of the enclosure that holds the managed systems, the I/O units, and the bulk power assemblies (BPAs).

I/O Units
The **I/O Units** tab displays all of the I/O units that are contained in the frame, their cage numbers, and their assigned managed systems. A cage is a division of the enclosure that holds the managed systems, the I/O units, and the BPAs. If the System column displays **Not owned**, the corresponding I/O unit is not assigned to a managed system.

Operations

Perform tasks on managed frames.

Initialize Frames

Learn how to initialize managed frames.

This operation task is available when one or more frames are selected. It first powers on the unowned I/O units within the selected managed frames, then power on the managed systems within the selected managed frames. The complete initialization process might take several minutes to complete.

Note: Managed systems that are already powered on are not affected and are not powered off and back on again.

Initialize All Frames

Initialize all of your frames.

About this task

This operation task is available when no managed frame is selected and the **Frames** tab on the navigation area is highlighted. It first powers on unowned I/O units within each managed frame, then power on managed systems within each managed frame.

Note: Frames are already powered on when they are connected to HMC. Initializing frames does not power on the frames.

Rebuild

Update frame information on the Hardware Management Console (HMC) interface.

Updating, or rebuilding, the frame acts much like a refresh of the frame information. Rebuilding the frame is useful when the system's state indicator in the Work pane of the HMC is shown as **Incomplete**. The **Incomplete** indicator signifies that the HMC cannot gather complete resource information from the managed system within the frame.

No other tasks can be performed on the HMC during this process, which can take several minutes.

Change Password

Change the Hardware Management Console (HMC) access password on the selected managed frame.

After the password is changed, you must update the HMC access password for all other HMCs from which you want to access this managed frame.

Enter the current password and then, enter a new password and verify it by entering it again.

Power On/Off IO Unit

Power off an IO unit by using the Hardware Management Console (HMC) interface.

Only units or slots that reside in a power domain can be powered off. The corresponding power on/off buttons are disabled for location codes that are not controllable by the HMC.

Configuration

You can use the **Configuration** tasks to configure your frame. You can also manage custom groups by using the **Configuration** task.

Manage Custom Groups

You can report status on a group basis to monitor your system in a way that you prefer.

You can also nest groups (a group that is contained within a group) to provide hierarchical or topology views.

One or more user-defined groups might already be defined on your HMC. Default groups are listed under **Custom Groups** node under **Server Management**. The default groups are **All Partitions** and **All Objects**. You can create others, delete the ones that were created, add to created groups, create groups by using the pattern match method, or delete from created groups by using the **Manage Custom Groups** task.

Use the online Help if you need additional information for working with groups.

Connections

You can use the **Connections** tasks to view the Hardware Management Console (HMC) connection status to frames or reset those connections.

Bulk Power Assembly (BPA) Status

Use the **Bulk Power Assembly (BPA) Status** task to view the state of the connection from the Hardware Management Console (HMC) to side A and side B of the bulk power assembly. The HMC operates normally with a connection to either side A or side B. However, for code update operations and some concurrent maintenance operations, the HMC needs connections to both sides.

The HMC displays the following information:

- IP address
- BPA Role
- Connection Status
- Connection Error code

If the status is not Connected, the Connection status might be one of the following conditions:

Starting/Unknown
One of the Bulk Power Assemblies (BPAs) contained in the frame is in the process of starting. The state of the other BPA cannot be determined.

Standby/Standby
Both of the BPAs contained in the frame are in the standby state. A BPA in the standby state is operating normally.

Standby/Starting
One of the BPAs contained in the frame is operating normally (in standby state). The other BPA is in the process of starting.

Standby/Not Available
One of the BPAs contained in the frame is operating normally (in the standby state), but the other BPA is not operating normally.

Pending frame number
A change to the frame number is in progress. No operations can be completed when the frame is in this state.

Failed Authentication
The HMC access password for the frame is not valid. Enter a valid password for the frame.

Pending Authentication - Password Updates Required
The frame access passwords are not set. You must set the required passwords for the frame to enable secure authentication and access control from the HMC.

No Connection
The HMC cannot connect to the frame.

Incomplete
The HMC failed to get all of the necessary information from the managed frame. The frame is not responding to requests for information.

Reset

Reset the connection between the Hardware Management Console (HMC) and the selected managed frame.

When you reset the connection with a managed frame, the connection is broken and then reconnected. Reset the connection with the managed frame if the managed frame is in a **No Connection** state and you verify that the network settings are correct on both the HMC and the managed frame.

Serviceability

Problem analysis on the Hardware Management Console (HMC) automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. You can view specific events for selected systems and add, remove, or exchange a Field Replaceable Unit (FRU). Use the **Serviceable Events Manager** task to view specific events for selected frames.

To open the serviceability tasks that are available for your frame, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Frames**.
2. Select the frame for which you want to manage serviceability tasks.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the serviceability task that you want to complete from the list.

Serviceable Events Manager

Problems on your managed frame are reported to the Hardware Management Console (HMC) as serviceable events. You can view the problem, manage problem data, call home the event to your service provider, or repair the problem.

To set the criteria for the serviceable events you want to view, complete the following steps:

1. From the menu pod, open **Serviceable Events Manager**.
2. Provide event criteria, error criteria, and FRU criteria.
3. Click **OK**.
4. If you do not want the results that are filtered, select **ALL**.

The **Serviceable Events Overview** window displays all of the events that match your criteria. The information that is displayed in the compact table view includes the following fields:

- Problem Number.
- PMH Number.
- Reference Code: Click **Reference** code to display a description of the problem reported and actions that can be taken to fix the problem.
- Status of the problem.
- Last reported time of the problem.
- Failing MTMS of the problem.

The full table view includes more detailed information, including reporting MTMS, first reported time, and serviceable event text.

Select a serviceable event and complete the following tasks:

- **View event details**: FRUs associated with this event and the descriptions.
- **Repair the event**: Start a guided repair procedure, if available.
- **Call home the event**: Report the event to your service provider.
- **Manage event problem data**: View, call home, or offload to media data and logs associated with this event.
- **Close the event**: After the problem is solved, add comments and close the event.

Use the online Help if you need additional information on managing serviceable events.

Hardware

These tasks are used to add, exchange, or remove hardware from the managed frame. From the hardware tasks, you can display a list of installed FRUs or enclosures and their locations. Select a FRU or an enclosure and start a step-by-step procedure to add, exchange, or remove the unit.

Add FRU

Use the **Add FRU** task to locate and add a FRU.

To add a FRU, complete the following steps:

1. From the **FRU** menu, select an enclosure type.
2. Select a FRU type.
3. Click **Next**.
4. Select a location code.
5. Add the selected enclosure location to Pending Actions by clicking **Add**.
6. Begin adding the selected FRU type to the enclosure locations identified in Pending Actions by clicking **Launch Procedure**.
7. When you complete the FRU installation process, click **Finish**.

Add Enclosure

Use the **Add Enclosure** task to locate and add an enclosure.

To add an enclosure, complete the following steps:

1. Select an enclosure type, then click **Add** to add the selected enclosure type's location code to **Pending Actions**.
2. To begin adding the enclosures that are identified in **Pending Actions** to the selected system, click **Launch Procedure**.
3. When you complete the enclosure installation process, click **Finish**.

Exchange FRU

Exchange one FRU with another FRU.

To exchange a FRU, complete the following steps:

1. Select an installed enclosure type.
2. Select a FRU type.
3. Click **Next**.
4. Select a location code for a specific FRU.
5. Click **Add**.
6. Select **Launch Procedure**.

   Note: This procedure identifies the resources that are impacted by the **Exchange FRU** task, including any resources that are in use by partitions. Workloads that are running on partitions might be impacted if redundancy is not configured. Follow the on-screen instructions to complete the exchange.
7. When you complete the installation, click **Finish**.

Exchange Enclosure

Exchange one enclosure for another enclosure.

To exchange an enclosure, complete the following steps:

1. Select an installed enclosure, then click **Add** to add the selected enclosure's location code to **Pending Actions**.
2. Begin replacing the enclosures that are identified in **Pending Actions** in the selected system by clicking **Launch Procedure**.

3. When you complete the enclosure replacement process, click **Finish**.

Remove FRU

Remove a FRU from your managed system.

To remove a FRU, complete the following steps:

1. Select an enclosure from the menu.
2. Select a FRU type from the displayed list of FRU types for this enclosure.
3. Click **Next**.
4. Select a location code for a specific FRU.
5. Click **Add**.
6. Select **Launch Procedure**.
7. When you complete the removal procedure, click **Finish**.

Remove Enclosure

Remove an enclosure that is identified by the Hardware Management Console (HMC).

To remove an enclosure, complete the following steps:

1. Select an enclosure type, then click **Add**.
2. Click **Launch Procedure**.
3. When you complete the enclosure removal process, click **Finish**.

Manage Groups

The **All groups** view provides a mechanism for you to group system resources together in a single view.

Groups may be nested to create customized system resources view.

You can view all the groups that are created by the users of the management console, including cumulative state information for system resources in a group. A custom group can consist of any systems, partitions, and Virtual I/O Servers that are managed by the management console.

To create a new group, complete the following steps:

1. Click **Create group** on the toolbar.
2. In the **Create group** window, specify a group name and description for the group. You can also tag a color to the group that you want to create.
3. Select one or more resources (for example: servers, partitions, or frames) that you want to include in the group that you want to work with.
4. Click **OK** to save the changes and to close the window.

You can edit an existing group to add or remove the resources from the group.

Note: When the last member (resources) of the group is removed, a message is displayed to confirm whether you want to delete the group. Click **Cancel** to retain the group in the **All groups** view.

Power Enterprise Pools

Systems Management for Power Enterprise Pool displays Power Enterprise Pool tasks that you can perform.

You can perform the following operations by using the Power Enterprise Pool offering:

- Add processors or memory to a server.
- Remove processors or memory from a server.
- Update the pool configuration.
- Add a server to the pool.

- Remove an existing server from the pool.
- Add processors or memory to the pool.
- View the following Power Enterprise Pool information:
  - Pool membership information
  - Pool resource information
  - Pool compliance information
  - Pool history log

HMC Management tasks

Learn about the tasks that are available on the Hardware Management Console (HMC) under **HMC Management**.

To open these tasks, see “HMC tasks, user roles, IDs, and associated commands” on page 9.

Note: Depending on the task roles that are assigned to your user ID, you might not have access to all the tasks. See Table 3 on page 9 for a listing of the tasks and the user roles that are allowed to access them.

Launch Guided Setup Wizard

This task uses a wizard to set up your system and HMC.

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Launch Guided Setup Wizard**.
3. From the **Launch Guided Setup Wizard - Welcome** window it is recommended that you have certain prerequisites on hand. Click **Prerequisites** in the **Launch Guided Setup Wizard - Welcome** window for the information. When you have completed that, this wizard takes you through the following tasks required to set up your system and HMC. As you complete each task, click **Next** to proceed.
   a. Change HMC Date and Time
   b. Change HMC passwords
   c. Create additional HMC users
   d. Configure HMC Network Settings (This task cannot be performed if you are accessing the **Launch Guided Setup Wizard** remotely.)
   e. Specify contact information
   f. Configure connectivity information
   g. Authorize users to use the Electronic Service Agent software tool and configure notification of problem events.
4. Click **Finish** when you have completed all the tasks in the wizard.

View Network Topology

This task allows you to view and ping the connectivity between various network nodes within the Hardware Management Console (HMC).

To view the network topology, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **View Network Topology**.
3. From the **View Network Topology** window, you can ping current and saved nodes.
4. Click **Close** when you have completed this task.

Use the online Help if you need additional information about viewing the network topology.

Test Network Connectivity

This task allows you to view network diagnostic information about the network protocols for the Hardware Management Console (HMC).

To test the network connectivity, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Test Network Connectivity**.
3. From the **Test Network Connectivity** window, you can work with the following tabs:

   Ping
   You can ping the TCP/IP address or name.

   Interfaces
   Displays the statistics for the network interfaces that are currently configured. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Ethernet Settings
   Displays the settings for the Ethernet cards that are currently configured. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Address
   Display the TCP/IP addresses for the configured network interfaces. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Routes
   Displays the Kernel IP and IPv6 routing tables and corresponding network interfaces. To update the information that is currently displayed with the most recent information, click **Refresh**.

   ARP
   Displays the contents of the Address Resolution Protocol (ARP) connections. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Sockets
   Displays information about TCP/IP sockets. To update the information that is currently displayed with the most recent information, click **Refresh**.

   TCP
   Displays information about Transmission Control Protocol (TCP) connections. To update the information that is currently displayed with the most recent information, click **Refresh**.

   IP Tables
   Displays information (in table format) about the Internet Protocol (IP) packet filter rules. To update the information that is currently displayed with the most recent information, click **Refresh**.

   UDP
   Displays information about User Datagram Protocol (UDP) statistics. To update the information that is currently displayed with the most recent information, click **Refresh**.
4. Click **Cancel** when you have completed this task.

Use the online Help if you need additional information about testing the network connectivity.

Change Network Settings

This task allows you to view the current network information for the HMC and to change network settings.

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change Network Settings**.
3. From the **Change Network Settings** window, you can work with the following tabs:

   Identification
   Contains the host name, domain name, and console description of the HMC.

   LAN Adapters
   A summarized list of all (visible) Local Area Network (LAN) adapters. You can select any of these and click **Details...** to open a window allowing you to change addressing, routing, other LAN adapter characteristics, and firewall settings.

   Bond LAN Adapters
   Create or delete a Bond LAN adapter. A Bond LAN adapter combines two Ethernet interfaces into a single logical link. To change the settings of the Bond LAN adapter, select a Bond LAN adapter and click **Edit**. You can change the IP address, IP network mask, gateway, and the firewall settings of the Bond LAN adapter.

   Name Services
   Specify the DNS and domain suffix values for configuring the console network settings.

   Routing
   Specify the routing information and default gateway information for configuring the console network settings.

   The **Gateway address** is the route to all networks. The default gateway address (if defined) informs this HMC where to send data if the target station does not reside on the same subnet as the source. If your machine can reach all stations on the same subnet (usually a building or a sector within a building), but cannot communicate outside the area, it is usually because of an incorrectly configured default gateway.

   You can assign a specific LAN to be the **Gateway device** or you can choose “any.”

   You can select **Enable ’routed’** to start the routed daemon, which allows it to run and allows any routing information to be exported from the HMC.
4. Click **OK** when you have completed this task.

Note: Depending on the type of change that you make, the network or console automatically restarts or the console automatically reboots.

Use the online Help to get additional information for customizing the network settings.

Change Performance Monitoring Settings

The Performance and Capacity Monitor tool collects allocation and usage data for virtualized server resources. It displays data in the form of graphs and tables, which are viewable from the Performance and Capacity Monitor home page.

The Performance and Capacity Monitor gathers data and provides capacity reporting and performance monitoring. This information can help you to determine the available capacity and whether your resources might be overextended or underused. In addition, your interpretation of the graphs and tables might be useful for capacity planning and troubleshooting. For more information about The Performance and Capacity Monitor tool, see Using the Performance and Capacity Monitor.

The Performance and Capacity Monitor captures data only from the servers for which you choose to enable data collection.

To enable data collection, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change Performance Monitoring Settings**.
3. Specify the number of days for which you want to store performance data by typing in a number 1 - 366. Alternatively, you can click the up or down arrows next to **Number of days to store performance data** under **Performance Data Storage**.

   Note: By default, the HMC stores data for 180 days. However, you can specify the maximum number of days that the HMC stores data to 366 days.
4. Click the toggle switch in the **Collection** column next to the name of the server for which you want to collect data. Alternatively, you can click **All On** to enable data collection for all of the servers in your environment that the HMC manages.

   Note: You might be prevented from collecting data from all of the servers in your environment because storage space is limited. The HMC prohibits you from enabling data collection from more servers when the HMC determines that it might run out of estimated storage space.
5. Click **OK** to apply the changes and close the window. You can now review the collected data when you access the Performance and Capacity Monitor home page.

Change Date and Time

Change the time and date of the battery-operated Hardware Management Console (HMC) clock and add or remove time servers for the Network Time Protocol (NTP) service.

Use this task in the following situations:

- If the battery is replaced in the HMC.
- If your system is physically moved to a different time zone.

Note: The time setting adjusts automatically for Daylight Saving Time in the time zone you select.

To change the date and time, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change Date and Time**.
3. Click the **Customize Console Date and Time** tab.
4. Enter the date and time information.
5. Click **OK**.

To change the time server information, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change Date and Time**.
3. Click the **NTP Configuration** tab.
4. Provide the appropriate information for the time server.
5. Click **OK**.

If you need additional information for changing the date and time of the HMC or for adding or removing time servers for the Network Time Protocol (NTP) service, use the online Help.

Change Language and Locale

This task sets the language and location for the HMC. After you select a language, you can select a locale associated with that language.

The language and locale settings determine the language, the character set, and other settings specific to the country or region (such as formats for date, time, numbers, and monetary units). Changes made in the **Change Language and Locale** window affect only the language and locale for the HMC itself. If you access the HMC remotely, the language and locale settings on your browser determine the settings that the browser uses to display the HMC interface.

To change the language and locale on the HMC:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Change Language and Locale**.
3. From the **Change Language and Locale** window, choose the applicable language and locale.
4. Click **OK** to apply the change.

Use the online Help if you need additional information for changing the language and locale of the HMC.

Create Welcome Text

Create and display a welcome message or display a warning message that appears before users log on to the Hardware Management Console (HMC).

The text that you enter in the message input area for this task appears on the **Welcome** window after you initially access the console. You can use this text to notify users about certain corporate policies or security restrictions that apply to the system.

To create a welcome text, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Create Welcome Text**.
3. Enter the welcome text that you want to display in the text box.

   Note: A maximum of 8192 characters is allowed.
4. Click **OK**.

For more information about this task, use the online Help.

Console Default Settings

You can modify the default console settings on the Hardware Management Console (HMC).

You can also modify the number of days for which a certificate is valid.

Note: The certificate can be valid for maximum of 3650 days.

To modify the console default settings, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.
2. In the content pane, click **Console Default Settings**.
3. In the **Console Default Settings** window, you can specify the number of days for which the certificate will be valid and you can also configure the time out settings for a HMC session. If you are using HMC

9.1.940, or later, you can specify the maximum number of log in attempts to the HMC graphical user interface (GUI). You can enter a value in the range 3 - 50.

4. When you complete the task, click **OK**.

Use the online Help if you need additional information about this task.

Shut Down or Restart

This task enables you to shut down (power off the console) or to restart the console.

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Shut Down or Restart**.
3. From the **Shut Down or Restart** window, you can:
   - Select **Restart the HMC** to automatically restart the HMC once the shut down has occurred.
   - Do not select **Restart the HMC** if you do not want to automatically restart the HMC.
4. Click **OK** to proceed with the shut down, otherwise click **Cancel** to exit the task.

Use the online Help if you need additional information about shutting down or restarting the HMC.

Schedule Operations

Create a schedule for certain operations to be performed on the Hardware Management Console (HMC) itself without operator assistance.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule a backup of important HMC information to DVD to occur once, or set up a repeating schedule.

The **Scheduled Operations** task displays the following information for each operation:

- The processor that is the object of the operation.
- The scheduled date.
- The scheduled time.
- The operation.
- The number of remaining repetitions.

From the **Scheduled Operations** window you can:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

An operation can be scheduled to occur one time or it can be scheduled to be repeated. You are required to provide the time and date that you want the operation to occur. If the operation is scheduled to be repeated, you are asked to select:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)

- The total number of repetitions. (required)

The operation that can be scheduled for the HMC is:

Backup Critical Console Data
Schedules an operation to back up the critical console hard disk information for the HMC.

To schedule operations on the HMC, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Schedule Operations**.
3. From the **Schedule Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, point to **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
4. To return to the HMC workplace, point to **Options** and then click **Exit**.

Use the online Help to get additional information for scheduling an operation.

View Licenses

View the Licensed Internal Code that you agreed to for this Hardware Management Console (HMC).

You can view licenses at any time. To view licenses, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **View Licenses**.
3. Click any of the license links to view more information.

   Note: This list does not include programs and code that is provided under separate license agreements.
4. Click **OK**.

Update the Hardware Management Console

Learn how to update the internal code of the Hardware Management Console (HMC) and view system information and system readiness.

To update the HMC, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Management**.

2. In the content pane, click **Update the Hardware Management Console**. The **Install HMC Corrective Service Wizard** opens.
3. Click **Next** to start the update process.
4. Follow the steps in the wizard to complete the update operation.
5. Click **Finish** when you have completed this task.

Use the online Help if you need additional information about updating the Hardware Management Console.

Format Media

This task formats a diskette or USB 2.0 Flash Drive Memory Key.

You can format a diskette by supplying a user-specified label.

To format a diskette or USB 2.0 Flash Drive Memory Key, complete the following steps:

1. In the navigation area, click the **HMC Management** icon, and then select **Console Managment**.
2. In the content pane, click **Format Media**.
3. From the **Format Media** window, select the type of media you want to format, then click **OK**.
4. Make sure that your media is correctly inserted, then click **Format**. The **Format Media** progress window is displayed. When the media is formatted, the **Format Media Completed** window is displayed.
5. Click **OK** and then click **Close** to end the task.

Use the online Help if you need additional information about this task.

Backup Management Console Data

This task backs up (or archives) the data that is stored on your Hardware Management Console (HMC) hard disk that is critical to support HMC operations.

Back up the HMC data after changes are made to the HMC or information that is associated with logical partitions.

The HMC data that is stored on the HMC hard disk drive can be saved to a DVD-RAM on a local system, a remote system that is mounted to the HMC file system (such as NFS), or sent to a remote site by using File Transfer Protocol (FTP).

By using the HMC, you can back up all important data, such as the following data:

- User-preference files
- User information
- HMC platform-configuration files
- HMC log files
- HMC updates through Install Corrective Service.

Note: Use the archived data only along with a reinstallation of the HMC from the product CDs.

To back up the HMC critical data, complete the following steps:

1. In the navigation area, click the **HMC Management** icon, and then select **Console Managment**.
2. In the content pane, click **Backup Management Console Data**.

3. From the **Backup Management Console Data** window, choose the archive option that you want to complete.
4. Click **Next**, then follow the appropriate instructions that are associated with the option you chose.
5. Click **OK** to continue with the backup process.

Use the online Help if you need additional information for backing up the HMC data.

Notes:

- For HMC model 7063-CR1, DVD media is not supported.
- If you are using HMC Version 9.1.940, or later, you can specify a name for the generated backup file. If the backup file exists on the server, select the **Replace file** to replace the contents of the existing file that has the same name.

Restore Management Console Data

This task is used to select a remote repository for restoring critical backup data for the HMC.

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Restore Management Console Data**.
3. From the **Restore Management Console Data** window, click **Restore from a remote Network File System (NFS) server**, **Restore from a remote File Transfer Protocol (FTP) server**, **Restore from a remote Secure Shell File Transfer Protocol (SFTP) server**, or **Restore from a remote removable media**.
4. Click **Next** to proceed or **Cancel** to exit the task without making any changes.

Use the online Help if you need additional information about restoring critical backup data for this HMC.

Save Upgrade Data

This task uses a wizard to save upgrade data to selected media. This data consists of files that were created or customized while running the current software level. Saving this data to selected media is performed prior to an HMC software upgrade.

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Save Upgrade Data**.
3. From the **Save Upgrade Data** window, this wizard takes you through the steps required for saving your data. Select the type of media you want to save your data to, then click **Next** to proceed through the task windows.
4. Click **Finish** when you have completed the task.

Use the online Help if you need additional information for saving upgrade data.

Manage Data Replication

This task enables or disables customized data replication. Customized data replication allows another HMC to obtain customized console data from or send data to this HMC.

The following types of data can be configured:

- Customer information data

- Administrator information (such as customer name, address, and telephone number)
  - System information (such as administrator name, address, and telephone of your system)
  - Account information (such as customer number, enterprise number, and sales branch office)
- Group data
  - All user-defined group definitions
- Modem configuration data
  - Configure modem for remote support
- Outbound connectivity data
  - Configure local modem to RSF
  - Enable an internet connection
  - Configure to an external time source

Note: Customizable console data is accepted from other HMCs only after specific HMCs and their associated allowable customizable data types have been configured.

To manage data replication, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Manage Data Replication**.
3. From the **Manage Data Replication** window, choose the appropriate option that you want to perform.

Use the online Help to get additional information for enabling or disabling customizable data replication.

Templates and OS Images

System templates contain configuration details for resources such as system properties, shared processor pools, reserved storage pool, shared memory pool, Host Ethernet adapters, and SR-IOV adapters. Many of the system settings that you previously configured by using separate tasks are available in the Deploy System from Template wizard. For example, you can configure the Virtual I/O Servers, virtual network bridges, and virtual storage settings when you use the wizard to deploy a system from a system template.

The template library includes predefined system templates, which contain configuration settings based on common usage scenarios. Predefined system templates are available for your immediate use. You can view, modify, deploy, copy, import, export, or delete templates that are available in the template library.

You can also create custom system templates that contain configuration settings that are specific to your environment. You can create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration.

To access the Template Library, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Templates and OS Images**.
2. From the **Templates and OS Images** window, you can access:
   - **System**
   - **Partition**
   - **OS and VIOS Images**
3. When you complete this task, click **Close**.

System Templates

System templates contain configuration information about resources such shared processor pools, reserved storage pool, shared memory pool, physical I/O adapters, Host Ethernet adapters, single root I/O virtualization (SRIOV) adapters, Virtual I/O Server (VIOS), virtual networks, and virtual storage.

You can create custom system templates that contain configuration settings that are specific to your environment. You can also create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration. Click the template name to see the details about the template. Select a system template from the list to view, edit, copy, delete, deploy, or export a template.

Use the online Help if you need additional information on system templates.

Partition Templates

Partition templates contain details about partition resources, such as physical adapters, virtual networks, and storage configuration.

You can create custom partition templates that contain configuration settings that are specific to your environment. You can also create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration. Click the template name to see the details about the template. Select a partition template from the list to view, edit, copy, delete, deploy, or export a template.

Notes:

- If you are using HMC Version 9.1.940, or later, and if you are using a non-captured template to create a logical partition, you can configure an SR-IOV logical port that can be migrated. Select **migratable** in the **Edit** menu of the partition template. You can migrate the logical partition by using the SR-IOV logical port by creating a backup device and associate the SR-IOV logical port to the logical partition. The backup device can either be a virtual Ethernet adapter or a virtual Network Interface Controller (NIC) adapter.
- When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Use the online Help if you need additional information on partition templates.

VIOS Images

Define VIOS images and installation resources for the operating environment that the Hardware Management Console (HMC) can access and use.

You can access the following tasks:

Manage Virtual I/O Server Image Repository

As of HMC version 7.7, or later, you can store the Virtual I/O Server (VIOS) images from a DVD, a saved image, or a Network Installation Management (NIM) server on the HMC. The stored VIOS images can be used for VIOS installation. You must be an HMC super administrator (hmcsuperadmin) to install the VIOS image.

About this task

To manage or to import the VIOS image repository, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Templates and OS Images**.
2. From the **Templates and OS Images** window, select the **OS and VIOS Images** tab, and then click **Manage Virtual I/O Server Image Repository**.
3. In the Virtual I/O Server Image Repository window, click **Import New Virtual I/O Server Image**.
4. In the Import New Virtual I/O Server Image window, choose to import the VIOS images from a DVD or from a file system.

   - To import VIOS images from a DVD to the HMC, complete the following steps:

     a. In the Import Virtual I/O Server Image window, select **Management console DVD**.

     b. In the **Name** field, enter the VIOS image name that you want to import from the DVD.

     c. Click **OK**.

   - To import VIOS images from a Network File System (NFS), File Transfer Protocol (FTP), or Secure Shell File Transfer Protocol (SFTP), complete the following steps:

     a. In the Import Virtual I/O Server Image window, select **File System**.

     b. Select **Remote NFS Server**, **Remote FTP Server**, or **Remote SFTP Server**.

     c. Enter the required details and click **OK**.

Manage Virtual I/O Server Backups

With HMC version 9.2.950, or later, you can manage the I/O configuration of Virtual I/O Servers and manage the backup of the VIOS image on the management console.

About this task

To manage the backup or restore operation of the I/O configuration of the VIOS and to manage the VIOS image, complete the following steps:

Procedure

1. In the navigation area, click the **HMC Management** icon , and then select **Templates and OS Images**.
2. From the **Templates and OS Images** window, select the **VIOS Images** tab, and then click **Manage Virtual I/O Server Backups**.
3. In the Manage Virtual I/O Server Backups window, select the **Virtual I/O Server Configuration Backup** tab. A table is displayed that lists all the backup files of the VIOS configuration that is taken by the HMC. Additionally, you can view the time at which the configuration file was last edited.

   a) To take the backup of the input/output configuration of a VIOS, click **Backup I/O configuration**. In the Backup I/O configuration window, select the managed system and the VIOS for which the back up is created, and then specify a name for the backup file.

   The name you specify must consist of 1 - 40 characters including file extension **`.tar.gz`**. You can use the characters A - Z and a - z, the numbers of 0 - 9, the dot (.), the dash (-) and the underscore (_) characters.

   b) To rename an existing backup file that is stored in the HMC, select a configuration file from the table and click **Action** > **Rename**.

   c) To restore the VIOS input/output configuration, select a backup file which contains the I/O configuration of the VIOS that you want to restore, and click **Action** > **Restore**.

4. In the Manage Virtual I/O Server Backups window, click the **Virtual I/O Server Backup** tab. A table is displayed that list all the VIOS image backup that are taken in the HMC. Additionally, you can also view the name and size of the VIOS image, the time when the VIOS image file was last edited, the managed system and the VIOS from which the image was captured.
   a) To take the backup of the VIOS image, click **Create Backup**. In the Create Backup window, select the managed system and the VIOS for which the back up is created, and then specify a name for the backup file.

      The name you specify must consist of 1 - 40 characters including file extension **`.tar`**. You can use the characters A - Z and a - z, the numbers of 0 - 9, the dot (.), the dash (-) and the underscore (_) characters.
   b) To rename an existing VIOS image backup file that is stored in the HMC, select a backup file from the table and click **Action** > **Rename**.
   c) To remove a VIOS image backup file from the HMC, select a backup file which contains the VIOS configuration that you want to remove from the table, and click **Action** > **Remove**.
5. Click **OK**.

All System Plans

A system plan is a specification of the logical partition configuration of a single managed system.

The table lists all the system plans that can be used to configure a managed system. You can create your own system plan or import an existing system plan.

Create System Plan

You can create a new system plan for a system that this Hardware Management Console (HMC) manages. The new system plan contains specifications for the logical partitions and partition profiles of the managed system that you used to create the plan.

1. Click **Create**.
2. Select a managed system from the available list and complete the **System plan name** and **Plan description** fields.
3. Check any options that you want.
4. Click **Create**.

Import System Plan

You can import a system plan file to the Hardware Management Console (HMC). The new system plan contains specifications for the logical partitions and partition profiles of the managed system that you used to create the plan.

1. Click **Import**.
2. Select a source to import the system plan file to the HMC.
3. Click **Import**.

Export System Plan

You can export a system plan file from the Hardware Management Console (HMC).

1. Select the system plan from the list and click **Actions** > **Export**.
2. Select a source to export the system plan file to the HMC.
3. Click **Export**.

Deploy System Plan

You can deploy a system plan file to one or more systems that the HMC manages. The managed system that you deploy the system plan on must have hardware that is identical to the hardware in the system plan.

1. Select the system plan from the list and click **Actions** > **Deploy**.
2. Follow the instructions on the **Deploy System Plan** wizard.

Delete System Plan

You can delete a system plan file from the Hardware Management Console (HMC).

1. Select the system plan from the list and click **Actions** > **Delete**.

Refresh

You can refresh the table to see any recent changes to the available system plans.

1. Click **Refresh** to update the table with the latest data.

Use the online Help if you need additional information about this task.

Users and Security tasks

The tasks that are available on the HMC for the **Users and Security** tasks are described.

Note: Depending on the task roles assigned to your user ID you may not have access to all the tasks. See "HMC tasks, user roles, IDs, and associated commands" on page 9 for a listing of the tasks and the user roles allowed to access them.

Change User Password

This task allows you to change your existing password that is used for logging on to the Hardware Management Console (HMC). A password verifies your user ID and your authority to log in to the console.

To change your password, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. In the content pane, click **Change User Password**.
3. From the **Change User Password** window, specify your current password, specify a new password that you want to use, and re-specify the new password for confirmation in the fields provided.

   Note: The new password that you specify must have atleast eight characters.
4. Click **OK** to proceed with the changes.

Use the online Help if you need additional information for changing your password.

Manage User Profiles and Access

Manage your system users that log on to the HMC. A user profile is a combination of a user ID, server authentication method, permissions, and a text description. Permissions represent the authority levels assigned to the user profile for the objects the user has permission to access.

Users can be authenticated using local authentication on the HMC, by using Kerberos remote authentication, or by using LDAP authentication. For more information on setting up Kerberos authentication on the HMC, see “Manage KDC” on page 92. For more information about LDAP authentication, see “Manage LDAP” on page 92.

For security reasons, remotely authenticated Kerberos or LDAP users cannot lock the local console.

If you are using local authentication, the user ID and password are used to verify a user’s authorization to log on the HMC. The user ID must start with an alphabetic character and consist of 1 to 32 characters. The password has the following rules:

- Must begin with an alphanumeric character.
- Must contain at least 8 characters, however, this limit can be changed by your system administrator.
- The characters must be standard 7-bit ASCII characters.
- Valid characters to use for the password can be: A-Z, a-z, 0-9 and special characters (~ ! @ # $ % ^ & * ( ) _ + - = { } [ ] \ : " ; ’).

If you are using Kerberos authentication, specify a Kerberos remote user ID.

If you select LDAP authentication, no additional information is required.

The user profile includes managed resource roles and task roles that are assigned to the user. The *managed resource roles* assign permissions for a managed object or group of objects and the *task roles* define the access level for a user to perform on a managed object or group of objects. You can choose from a list of available default managed resource roles, task roles, or customized roles that are created by using the **Manage Task and Resource Roles** task.

See “HMC tasks, user roles, IDs, and associated commands” on page 9 for a listing of all the HMC tasks and the predefined default user IDs that can perform each task.

The default managed resource roles include:

- All System Resources

The default task roles include:

- hmcservicerep (Service Representative)
- hmcviewer (Viewer)
- hmcoperator (Operator)
- hmcpe (Product Engineer)
- hmcsuperadmin (Super Administrator).

To add or customize a user profile, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. In the content pane, click **Manage User Profiles and Access**.
3. Complete one of the following steps:
   - From the **User Profiles** window, if you are creating a new user ID, point to **User** on the menu bar and when its menu is displayed, click **Add**. The **Add User** window is displayed.
   - From the **User Profiles** window, if you are creating a user ID with the same attributes as an existing profile, point to **User** on the menu bar and when its menu is displayed, click **Copy**. The **Copy User** window is displayed.

     Note: Some user profiles are predefined, such as a default ID, and those permissions cannot be changed. However, you can copy a default user profile, such as operator, and then modify the resulting new user profile. The newly defined user cannot have greater permissions than the original copied user profile.
   - From the **User Profiles** window, if you are deleting a user ID, point to **User** on the menu bar and when its menu is displayed, click **Remove**. The **Remove User** window is displayed.

- From the **User Profiles** window, if the user ID exists in the window, select the user ID from the list, and then point to **User** on the menu bar and when its menu is displayed, click **Modify**. The **Modify User** window is displayed.
  - To specify timeout and inactivity values, click **User Properties** from the **Modify User** window.

4. Complete or change the fields in the window, click **OK** when you are done.

Use the online Help if you need additional information for creating, modifying, copying, or removing a user profile and modifying timeout and inactivity values.

Adding, Copying, or Modifying User Profiles

Learn how to add, copy, or modify user profiles.

Users who remotely authenticate through Kerberos or Lightweight Directory Access Protocol (LDAP) must have their profiles set appropriately. You must set the user profile of each remotely authenticated Kerberos or LDAP user to use that type of authentication instead of local authentication. A user that is set to use Kerberos or LDAP remote authentication always uses that type of authentication, even when the user logs into the HMC locally.

Note: The use of Kerberos authentication requires configuration of a key distribution center (KDC) server by using the **KDC Configuration** task. The use of LDAP authentication requires configuration of an LDAP server by using the **LDAP Configuration** task. You do not need to set all users to use Kerberos or LDAP remote authentication. You can set some user profiles so that the users can use local authentication only.

From the Adding, Copying, or Modifying User Profiles window, you can modify the following attributes:

- **User ID**: Enter the user ID for the user profile you are creating or managing. The user name must start with an alphabetic character and consist of 1 to 32 characters.
- **Description**: Enter a meaningful description for your own records.
- **Password**: Enter the password for the user ID.
- **Confirm password**: Enter the password again for verification.
- **Password expires in (days)**: Specify the number of days a password is valid before it expires. This input field is available when **Enforce strict password rules** check box is selected.
- **Manage resource roles**: Displays the managed resource roles that are currently available. Select one or more managed resource roles to define access permissions for this user ID.
- **Task roles**: Displays the task roles that are currently available. Select one task role for this user ID.

Use the online Help if you need additional information about creating, modifying, copying, or removing a user profile and modifying timeout and inactivity values.

User Properties

Learn how to specify timeout and inactivity values for the selected user.

You can specify the amount of time for the following timeout and inactivity tasks:

Timeout Values

- **Session timeout minutes**: Specifies the number of minutes during a logon session that a user is prompted for identity verification. If a value other than zero is specified, the user is prompted after the specified time is reached to reenter their password. If a password is not reentered within the specified amount of time in the **Verify timeout minutes** field, the session is disconnected.
- **Verify timeout minutes**: Specifies the amount of time that is required for the user to reenter their password when prompted, if a value was specified in the **Session timeout minutes** field. If the password is not reentered within the specified time, the session is disconnected.
- **Idle timeout minutes**: Specifies the number of minutes the user's session can be idle. If the user does not interact with the session in the specified amount of time, the session is locked and the screen saver starts. Clicking anywhere on the screen prompts the user for identity verification.

- **Minimum time in days between password changes**: Specifies the minimum amount of time in days that must elapse between changes for the user's password.

Note: A note of zero in any of these fields indicates that there is no expiration of time and it is the default value. You can specify up to a maximum value of 525600 minutes (equivalent to one year).

Inactivity Values

- **Disable for inactivity in days**: Specifies the amount of time in days a user is temporarily disabled after the maximum number of days of inactivity is reached.
- **Never disable for inactivity**: Option to never disable a user's session due to inactivity.
- **Allow remote access via the web**: Option to enable remote web server access for the user you are managing.

Manage Users and Tasks

Display the logged on users and the tasks they are running.

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Manage Users and Tasks**.
3. In the Manage Users and Tasks window, the following information displays:
   - User you are logged in as
   - Time you logged in
   - Number of tasks running
   - Your access location
   - Information about tasks that are running:
     - Task ID
     - Task name
     - Targets (if any)
     - Session ID
4. Choose to log off or disconnect from a session that is currently running by selecting the session from the Users **Logged On** list, then click **Logoff** or **Disconnect**.

   Alternately, you can choose to switch to another task or end a task by selecting the task from the **Running Tasks** list, then click **Switch To** or **Terminate**.
5. When you have completed this task, click **Close**.

Manage Task and Resource Roles

Use this task to define and customize user roles.

Note: Predefined roles (default roles) cannot be modified.

A *user role* is a collection of authorizations. A user role can be created to define the set of tasks allowed for a given class of user (*task roles*) or it can be created to define the set of managed objects that are manageable for a user (*managed resource roles*). Once you have defined or customized the user roles you can use the **Manage User Profiles and Access** task to create new users with their own permissions.

If the automatic resource role update function is enabled on the Hardware Management Console (HMC) either through the command line interface or through the Rest API CLI runner job, the HMC user can automatically receive permission to the logical partition that is created. If the logical partition is deleted, the permission is automatically revoked.

The predefined managed resource roles include:

- All System Resources

The predefined task roles include:

- hmcservicerep (Service Representative)
- hmcviewer (Viewer)
- hmcoperator (Operator)
- hmcpe (Product Engineer)
- hmcsuperadmin (Super Administrator)

To customize managed resource roles or task roles:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Manage Task and Resource Roles**.
3. From the **Manage Task and Resource Roles** window, select either **Managed Resource Roles** or **Task Roles**.
4. To add a role, click **Edit** from the menu bar, then click **Add** to create a new role.

   or

   To copy, remove, or modify an existing role, select the object you want to customize, click **Edit** from the menu bar, then click **Copy**, **Remove**, or **Modify**.
5. Click **Exit** when you are have completed the task.

Use the online Help to get additional information for customizing managed resource roles and task roles.

Manage Certificates

Use this task to manage the certificates used on your HMC. It provides the capability of getting information on the certificates used on the console. This task allows you to create a new certificate for the console, change the property values of the certificate, and work with existing and archived certificates or signing certificates.

All remote browser access to the HMC must use Secure Sockets Layer (SSL) encryption. With SSL encryption required for all remote access to the HMC, a certificate is required to provide the keys for this encryption. The HMC provides a self-signed certificate that allows this encryption to occur.

To manage your certificates:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Manage Certificates**.
3. Use the menu bar from the **Manage Certificates** window for the actions you want to take with the certificates:
   - To create a new certificate for the console, click **Create**, then select **New Certificate**. Determine whether your certificate will be self-signed or signed by a Certificate Authority (CA), then click **OK**.
   - To modify the property values of the self-signed certificate, click **Selected**, then select **Modify**. Make the appropriate changes, then click **OK**.

  Note: If you have a certificate signed by a Certificate Authority (CA) that consists of a root certificate, intermediate certificate, and a client or leaf certificate, complete the following steps to upload the certificate to the HMC:

  - Open the CA signed certificate file by using a text-based editor and split the content of the file and save as three separate files. The first file is the client or leaf certificate, the second file is the intermediate certificate, and the third file is the root certificate.
  - Log in to the HMC to import the certificate. First upload the client certificate and click **Yes** for uploading more files. In the new window, upload the intermediate certificate and the root certificate.
  - Click **OK** to restart the console.

- To work with existing and archived certificates or signing certificates, click **Advanced**. Then you can choose the following options:
  - Delete existing certificates
  - Work with archived certificates
  - Import certificates
  - View issuer certificates

4. Click **Apply** for all changes to take effect.

Use the online Help if you need additional information for managing your certificates.

Manage Certificate Revocation List

Use this task to create, modify, delete, and import the certificate revocation list that is used on your Hardware Management Console (HMC).

All remote browsers that are accessing the HMC must use Secure Sockets Layer (SSL) encryption. A certificate is required to provide the keys for this encryption. The HMC provides a self-signed certificate that allows this encryption to occur.

To manage your certificate revocation list, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔒, and then select **Users and Roles**.
2. In the content pane, click **Manage Certificate Revocation List**.
3. Use the menu bar from the **Manage Certificate Revocation List** window for the actions you want to take with the certificates:
   - To create a new certificate revocation list for the console, click **Import**, then select **New CRL**. Determine whether your certification revocation list is imported from removable media on the console or from the file system on the system that is running the web browser.

     Note: If the list is from removable media, then the certificate revocation list file must be in the top directory on the media.
   - To modify a certificate revocation list on the console, select the certification revocation list from the table, and make appropriate changes, then click **Apply**.
   - To delete a certificate revocation list from the console, click **Selected**, then select **Delete CRL**. Select the certification revocation list, then click **OK**.
   - To work with existing and archived certificates or signing certificates, click **Advanced**.

Use the online Help if you need additional information for managing your certificate revocation list.

Manage LDAP

Configure your HMC so that it uses LDAP (Lightweight Directory Access Protocol) authentication.

Before you begin

Note: Before you configure the HMC so that it uses LDAP authentication, you must ensure that a working network connection exists between the HMC and the LDAP servers.

About this task

To configure your HMC so that it uses LDAP authentication, complete the following steps:

Procedure

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Systems and Console Security**.
2. In the content pane, click **Manage LDAP**. The **LDAP Server Definition** window opens.
3. Select **Enable LDAP**.
4. Define an LDAP server to use for authentication (for example, Microsoft Active Directory, Tivoli®, and Open LDAP).
5. Define the LDAP attribute that is used to identify the authenticated user. The default is **uid**, but you can use your own attributes. For Microsoft Active Directory, use **sAMAccountName** as the attribute.
6. Define the distinguished name tree, also known as the search base, for the LDAP server.
7. Click **OK**.

What to do next

If you want to use LDAP authentication, you must configure each remote user's profile so that it uses LDAP remote authentication instead of local authentication.

Manage KDC

View the key distribution center (KDC) servers that are used by this Hardware Management Console (HMC) for Kerberos remote authentication.

From this task, you can complete the following tasks:

- View existing KDC servers.
- Modify existing KDC server parameters that include realm, ticket lifetime, and clock skew.
- Add and configure a KDC server on the HMC.
- Remove a KDC server.
- Import a service key.
- Remove a service key.

Kerberos is a network authentication protocol that is designed to provide strong authentication for client/server applications by using secret-key cryptography.

Under Kerberos, a client (generally either a user or a service) sends a request for a ticket to the KDC. The KDC creates a ticket-granting ticket (TGT) for the client, encrypts it using the client's password as the key, and sends the encrypted TGT back to the client. The client then attempts to decrypt the TGT, by using its password. If the client successfully decrypts the TGT (for example, if the client gave the correct password), it keeps the decrypted TGT, which indicates proof of the client's identity.

The tickets have a time availability period. Kerberos requires the clocks of the involved hosts to be synchronized. If the HMC clock is not synchronized with the clock of KDC server, authentication fails.

A Kerberos realm is an administrative domain, site, or logical network that uses Kerberos remote authentication. Each realm uses a master Kerberos database that is stored on a KDC server and that contains information about the users and services for that realm. A realm might also have one or more slave KDC servers, which store read-only copies of the master Kerberos database for that realm.

To prevent KDC spoofing, the HMC can be configured to use a service key to authenticate to the KDC. Service key files are also known as keytabs. Kerberos verifies that the TGT requested was issued by the same KDC that issued the service key file for the HMC. Before you can import a service key file into an HMC, you must generate a service key for the host principal of the HMC client.

Note: For MIT Kerberos V5 *nix distributions, create a service key file by running the `kadmin` utility on a KDC and by using the `ktadd` command. Other Kerberos implementations might require a different process to create a service key.

You can import a service key file from one of these sources:

- Removable media that is mounted to the HMC, such as optical discs or USB Mass Storage devices. You must use this option locally at the HMC (not remotely), and you must mount the removable media to the HMC before you use this option.
- A remote site that uses secure FTP. You can import a service-key file from any remote site with SSH installed and running.

To use Kerberos remote authentication for this HMC, complete the following tasks:

- You must enable the Network Time Protocol (NTP) service on the HMC and set the HMC and the KDC servers to synchronize time with the same NTP server. You can enable the NTP service on the HMC by accessing the “Change Date and Time” on page 76 task from the **HMC Management** icon , and then selecting **Console Settings**.
- You must set the user profile of each remote user to use Kerberos remote authentication instead of local authentication. A user that is set to use Kerberos remote authentication always uses Kerberos remote authentication, even when the user logs on to the HMC locally.

  Note: You do not need to set all users to use Kerberos remote authentication. You can set some user profiles so that the users can use local authentication only.
- Use of a service key file is optional. Before you use a service key file, you must import it into the HMC. If a service key is installed on the HMC, realm names must be equivalent to the network domain name. The following example shows how to create the service key file on a Kerberos server by using the **kadmin.local** command, assuming the HMC hostname is `hmc1`, the DNS domain is `example.com`, and the Kerberos realm name is `EXAMPLE.COM`:

  –
  ```
  # kadmin_local kadmin.local: ktadd -k /etc/krb5.keytab host/
  hmc1.example.com@EXAMPLE.COM
  ```

  Using the Kerberos ktutil on the Kerberos server, verify the service key file contents. The output looks like the following example:

  –
  ```
  # ktutil

  ktutil: rkt /etc/krb5.keytab

  ktutil: l

  slot KVNO Principal

  ---- ----
  --------------------------------------------------------------------

  1 9 host/hmc1.example.com@EXAMPLE.COM

  2 9 host/hmc1.example.com@EXAMPLE.COM
  ```
- The HMC Kerberos configuration can be modified for SSH (Secure Shell) login without a password by using GSSAPI. For remote login without a password through Kerberos to an HMC, configure the HMC to

use a service key. After the configuration is completed, use `kinit -f principal` to obtain forwardable credentials on a remote Kerberos client machine. You can then enter the following command to log in to the HMC without having to enter a password: `$ ssh -o PreferredAuthentications=gssapi-with-mic user@host`.

To manage the KDC, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Manage KDC** window, select the appropriate task from the available options under the **Actions** menu.
4. When you complete the task, click **OK**.

Use the online Help if you need additional information for Managing KDC.

View KDC Server

Display existing key distribution center (KDC) servers on the Hardware Management Console (HMC).

To view existing KDC Servers on your HMC, click the **Users and Security** icon 🔓, and then select **Users and Roles**. In the content pane, click **Configure KDC**. If no servers exist and NTP has not yet been enabled, a warning panel message displays. Enable the NTP service on the HMC and configure a new KDC server as desired.

Modify KDC Server

Learn how to modify the key distribution center (KDC) on your Hardware Management Console (HMC).

To modify existing key distribution center (KDC) server parameters, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. Select a KDC Server.
4. Select a value to modify:
   - **Realm**. A realm is an authentication administrative domain. Normally, realms always appear in upper case letters. It is good practice to create a realm name that is the same as your DNS domain (in upper case letters). A user belongs to a realm if and only if the user shares a key with the authentication server of that realm. Realm names must be equivalent to the network domain name if a service key file is installed on the HMC.
   - **Ticket Lifetime**. Ticket lifetime sets the lifetime for credentials. The format is an integer number followed by one of **s** seconds, **m** minutes, **h** hours, or **d** days. Enter a Kerberos lifetime string such as *2d4h10m*.
   - **Clock skew**. Clock skew sets the maximum allowable amount of clock skew between the HMC and the KDC server before Kerberos considers messages invalid. The format is an integer number that represents number of seconds.
5. Click **OK**.

Add KDC server

Add a Key Distribution Center (KDC) server to this Hardware Management Console (HMC).

To add a new KDC server, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Actions** drop down list, select **Add KDC Server**.
4. Enter the host name or IP address of the KDC server.
5. Enter the KDC server realm.
6. Click **OK**.

Remove KDC server

Kerberos authentication on the Hardware Management Console (HMC) remains enabled until all key distribution center (KDC) servers are removed.

To remove a KDC server:

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. Select the KDC server from the list.
4. From the **Actions** drop down list, select **Remove KDC Server**.
5. Click **OK**.

Import Service Key

Before you can import a service key file into an Hardware Management Console (HMC), a service key file must first be created on the Kerberos server for the HMC host. The service key file contains the host principal of the HMC client, for example, `host/example.com@EXAMPLE.COM`. In addition to KDC Authentication, the host service key file is used to enable password-less SSH (Secure Shell) login using GSSAPI.

Note: For MIT Kerberos V5 *nix distributions, create a service key file by running the `kadmin` utility on a KDC and using the `ktadd` command. Other Kerberos implementations may require a different process to create a service key.

To import a service key, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Actions** drop down list, select **Import Service Key**.
4. Select from one of the following:
   - **Local** - The service key must be located on removable media currently mounted on the HMC. You must use this option locally at the HMC (not remotely), and you must mount the removable media to the HMC before using this option. Specify the full path of the service key file on the media.

- **Remote** - The service key must be located on a remote site available to the HMC via secure FTP. You can import a service key file from any remote site that has SSH (Secure Shell) installed and running. Specify the hostname of the site, a user ID and password for the site, and the full path of the service key file on the remote site.

5. Click **OK**.

Implementation of the service key file will not take effect until the HMC is rebooted.

Remove Service Key

Learn how to remove a service key from your Hardware Management Console (HMC).

To remove the service key from the HMC, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon , and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Actions** drop down list, select **Remove Service Key**.
4. Click **OK**.

You must reboot the HMC after removing the service key. Failure to reboot may cause login errors.

Manage MFA

Learn how to enable Multi-Factor Authentication (MFA) on the Hardware Management Console (HMC).

Notes:

1. Multi-Factor Authentication is disabled on the HMC by default.
2. For HMC GUI login, when MFA is enabled and the user is configured on the PowerSC MFA server, enter the Cache Token Credential (CTC) code in the password field.
3. For Secure Shell (SSH) login:

   When MFA is enabled, all users that login through SSH are prompted for a CTC code. If the user is configured on the PowerSC MFA server, then you can enter the CTC code at the prompt. If the user is not configured on the PowerSC MFA server, press `Enter` when prompted for CTC code, and then enter the password of the user at the prompt.

To enable Multi-Factor Authentication, complete the following steps:

1. In the navigation area, click the **Users and Security** icon , and then select **Systems and Console Security**.
2. In the content pane, click **Manage MFA**.
3. From the **Manage MFA** window, select the **Enable multi factor authentication** check box.
4. Enter the following information:
   - **Host name or IP address of the authentication server**
   - **Port of the authentication server**
5. Click **OK**.

Use the online Help if you need additional information about this task.

Enable Remote Command Execution

This task is used to enable remote command execution using the ssh facility.

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Command Execution**.
3. From the **Enable Remote Command Execution** window, select **Enable remote command execution using the ssh facility**.
4. Click **OK**.

Enable Remote Operation

This task is used to allow the HMC to be accessed at a remote workstation through a web browser.

To enable the HMC remote access:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Operation**.
3. Select **Enabled** from the Remote Operation drop-down list, then click **OK**. The HMC can be accessed from a remote workstation using a Web browser.

Use the online Help to get additional information for allowing remote access to the HMC.

Enable Remote Virtual Terminal

A Remote Virtual Terminal connection is a terminal connection to a logical partition from another remote HMC. Use this task to enable Remote Virtual Terminal access for remote clients.

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Virtual Terminal**.
3. From the **Enable Remote Virtual Terminal** window, you can enable this task by selecting Enable remote virtual terminal connections.
4. Click **OK** to activate your changes.

Use the online Help to get additional information for enabling a remote terminal connection.

Serviceability tasks

The tasks that are available on the HMC for the **Serviceability** tasks are described.

Note: Depending on the task roles assigned to your user ID you may not have access to all the tasks. See "HMC tasks, user roles, IDs, and associated commands" on page 9 for a listing of the tasks and the user roles allowed to access them.

Tasks Log

View all the tasks that are currently running or completed on the Hardware Management Console (HMC).

To view the tasks log, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Tasks Log**.
2. You can view the following tabs in the tasks log:
   - **Task name**: Displays the name of task.
   - **Status**: Displays the current state of the task (running or completed).
   - **Resource**: Displays the name of the resource.
   - **Resource type**: Displays the type of resource.
   - **Initiator**: Displays the name of the user that initiated the task.
   - **Start time**: Displays the time that the task was initiated.
   - **Duration**: Displays the amount of time that the task took to complete.

Use the online Help for additional information about viewing the tasks log.

Console Events Logs

View a record of system events occurring on the Hardware Management Console (HMC). System events are individual activities that indicate when processes occur, begin and end, succeed or fail.

To view console events legs, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Console Events Logs**.
2. Use the menu bar to change to a different time range, or to change how the events display in the summary. You can also use the table icons or the **Select Action** menu on the table toolbar to display different variations of the table.
3. When you are done viewing the events, select **View** on the menu bar, then click **Exit**.

Use the online Help for additional information about viewing HMC events.

Serviceable Events Manager

This task allows you to select the criteria for the set of serviceable events you want to view. When you finish selecting the criteria, you can view the serviceable events that match your specified criteria.

To set the criteria for the serviceable events you want to view, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Serviceable Events Manager**.
2. From the **Serviceable Events Manager** window, provide event criteria, error criteria, and FRU criteria.
3. Click **OK** when you have specified the criteria you want for the serviceable events you want to view.

Use the online Help if you need additional information managing events.

Events Manager for Call Home

Learn how to monitor and approve any data that is being transmitted from an HMC to IBM.

1. In the navigation area, click the **Serviceability** icon , and then select **Events Manager for Call Home**.
2. From the **Events Manager for Call Home** window, select **Manage Consoles** to manage the list of registered management consoles. You can use the **Event Criteria** to specify the approval state, status, and originating HMC to filter the list of events that are available for all registered management consoles. You can use the criteria to filter the view and select events to view details, view files, and complete call home operations.
3. Click **OK** to exit Events Manager for Call Home and to save the filter values.

Use the online Help if you need additional information about this task.

Create Serviceable Event

This task reports problems that occurred on your Hardware Management Console (HMC) to the service provider (for example, the mouse does not work) or lets you test problem reporting.

Submitting a problem is dependent upon whether you have customized this Hardware Management Console to use the Remote Support Facility (RSF) and if it is authorized to automatically call for service. If so, the problem information and service request is sent to the service provider automatically with a modem transmission.

To report a problem on your Hardware Management Console, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Create Serviceable Event**.
3. From the **Create Serviceable Event** window, select a problem type from the list displayed.
4. Enter a brief description of your problem in the **Problem Description** input field and then click **Request Service**.

To test problem reporting from the **Report a Problem** window:

1. Select **Test automatic problem reporting** and enter *This is just a test* in the **Problem Description** input field.
2. Click **Request Service**. The problems are reported to the service provider for the Hardware Management Console. Reporting a problem sends to the service provider the information you provide on **Report a Problem** window, and machine information that identifies the console.

Use the online Help if you need additional information for reporting a problem or testing if problem reporting works.

Manage Dumps

Learn how to manage the procedures for dumps of selected systems on your Hardware Management Console (HMC).

To manage a dump, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.

2. In the content pane, click **Manage Dumps**.
3. From the **Manage Dumps** window, select a dump and perform one of the following dump-related tasks:

   From **Selected** on the menu bar:

   - Copy the dump to media.
   - Copy the dump to a remote system.
   - Use the call home feature to transmit the dump to your service provider.
   - Delete a dump.

   From **Actions** on the menu bar:

   - Initiate a dump of the hardware and server firmware for the managed system.
   - Initiate a dump of the service processor.
   - Initiate a dump of the Bulk Power Control service processor.
   - Modify the dump capability parameters for a dump type.

   From **Status** on the menu bar, you can view the offload progress of the dump.
4. Click **OK** when you complete this task.

Use the online Help to get additional information for managing dumps.

Transmit Service Information

Transmit service information to your service provider immediately or schedule when to transmit service information for use for problem determination.

To schedule or transmit service information, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Transmit Service Information**.
3. In the content pane, click the **Schedule and Send Data** tab to schedule the service information.

   Note: You can also click the following tabs to select the data that you want to send and to configure FTP connections:

   - **Schedule and Send Data**: Transmit information to your service provider immediately or schedule the transmission.
   - **Configure FTP Connection**: Provide configuration data to allow the use of FTP to offload service information.
   - **Send Problem Reports**: Select the data that you want and the destination for the data.
4. Select the types of service information that you want to enable regular transmissions or to send immediately.

   - **Operational Test (Heartbeat) Information -- always enabled**: Send the Problem Event log file.
   - **Hardware Service Information (VPD)**: Send the Vital Product Data (VPD) for all managed systems that are attached to this HMC.
   - **Software Service Information**: Send the VPD for all software that is running on the partitions.
   - **Performance Management Information**: Gather and send the performance management information.
   - **Update Access Key Information**: Verifies and updates Access Key information.
5. Select the interval (in days) and the time to schedule repeating transmissions. To transmit the information immediately, click **Send Now**.
6. Click **OK**.

Use the online Help for additional information about scheduling service information.

Format Media

This task formats a diskette or USB 2.0 Flash Drive Memory Key.

You can format a diskette by supplying a user-specified label.

To format a diskette or USB 2.0 Flash Drive Memory Key, complete the following steps:

1. In the navigation area, click the **HMC Management** icon , and then select **Console Managment**.
2. In the content pane, click **Format Media**.
3. From the **Format Media** window, select the type of media you want to format, then click **OK**.
4. Make sure that your media is correctly inserted, then click **Format**. The **Format Media** progress window is displayed. When the media is formatted, the **Format Media Completed** window is displayed.
5. Click **OK** and then click **Close** to end the task.

Use the online Help if you need additional information about this task.

Electronic Service Agent Setup Wizard

Learn how to open the Electronic Service Agent Setup wizard using the Hardware Management Console (HMC) interface.

About this task

To open the Electronic Service Agent Setup wizard, complete the following steps:

Procedure

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the contents pane, select **Electronic Service Agent Setup Wizard**. The Electronic Service Agent wizard opens. Follow the instructions in the wizard to configure call-home tasks.

Authorize User

Request authorization for Electronic Service Agent. Electronic Service Agent associates your system with a user ID and allows access to system information through the Electronic Service Agent facility. This registration is also used by your operating system to automate service processes for your AIX or IBM i operation system.

To register a user ID, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Authorize User**.
3. Provide a user ID that is registered with the Electronic Service Agent. If you need a user ID, you can register at the IBM Registration website.
4. Click **OK**.

Use the online Help if you need additional information for registering a customer user ID with the eService website.

Enable Electronic Service Agent

This task allows you enable or disable the call-home state for managed systems.

Note: If Customizable Data Replication is **Enabled** on this HMC (using the **Manage Data Replication** task), the data specified in this task may change depending on automatic replication from other HMCs configured on your network. For more information on data replication, see "Manage Data Replication" on page 81.

By enabling the call-home state for a managed system this causes the console to automatically contact a service center when a serviceable event occurs. When a managed system is disabled, your service representative is not informed of serviceable events.

To manage call-home for the system(s):

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Enable Electronic Service Agent**.
3. From the **Enable Electronic Service Agent** window, select a system or systems you want to enable or disable the call-home state.
4. Click **OK** when you have completed the task.

Use the online Help if you need additional information for enabling the Electronic Service Agent.

Manage Outbound Connectivity

Customize the means for outbound connectivity for the Hardware Management Console (HMC) to use to connect to remote service.

Note: If Customizable Data Replication is **Enabled** on this HMC (using the **Manage Data Replication** task), the data specified in this task may change depending on automatic replication from other HMCs configured on your network. For more information on data replication, see "Manage Data Replication" on page 81.

You can configure this HMC to attempt connections through the local modem, Internet, Internet Virtual Private Network (VPN), or through a remote pass-through system. Remote service is a two-way communication between the HMC and the IBM Service Support System for the purpose of conducting automated service operations. The connection can only be initiated by the HMC. IBM Service Support System cannot and never attempts to initiate a connection to the HMC.

To customize your connectivity information, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Manage Outbound Connectivity**.
3. From the **Manage Outbound Connectivity** window select **Enable local server as call-home server** (a check mark appears) before proceeding with the task.

   Note: You must first **Accept** the terms described about the information you provided in this task.

   This allows the local HMC to connect to your service provider's remote support facility for call-home requests.
4. The dial information window displays the following tabs for providing input:
   - Local Modem
   - Internet
   - Internet VPN

- Pass-Through Systems

5. If you want to allow connectivity over a modem, use the **Local Modem** tab, then select **Allow local modem dialing for service** .
   a. If your location requires a prefix to be dialed in order to reach an outside line, click **Modern Configuration** and enter the **Dial prefix** in the **Customize Modem Settings** window required by your location. Click **OK** to accept the setting.
   b. Click **Add** from the **Local Modem** tab page to add a telephone number. When local modem dialing is allowed, there must be at least one telephone number configured.
6. If you want to allow connectivity over the Internet, use the **Internet** tab, then select **Allow an existing internet connection for service**.
7. If you want to configure the use of a VPN over an existing Internet connection to connect from the local HMC to your service provider's remote support facility, use the **Internet VPN** tab.
8. If you want to allow the HMC to use the pass-through systems as configured by the TCP/IP address or host name, use the **Pass-Through Systems** tab.
9. When you complete all the necessary fields, click **OK** to save your changes.

Use the online Help if you need additional information for customizing outbound connectivity information.

Manage Inbound Connectivity

Learn how to allow your service provider to temporarily access your local console, such as the Hardware Management Console (HMC), or the partitions of a managed system.

To manage inbound connectivity, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Manage Inbound Connectivity**.
3. From the **Manage Inbound Connectivity** settings window, you can perform the following tasks:
   - Use the **Remote Service** tab to provide the information necessary to start an attended remote service session.
   - Use the **Call Answer** tab to provide the information necessary to accept incoming calls from your service provider to start an unattended remote service session.
4. Click **OK** to proceed with your selections.

Use the online Help if you need additional information about this task.

Manage Customer Information

This task enables you to customize the customer information for the Hardware Management Console (HMC).

Note: If Customizable Data Replication is *Enabled* on this HMC (using the **Manage Data Replication** task), the data specified in this task may change depending on automatic replication from other HMCs configured on your network. For more information on data replication, see “Manage Data Replication” on page 81.

The **Manage Customer Information** window displays the following tabs for providing input:

- Administrator
- System
- Account

To customize your customer information, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Manage Customer Information**.
3. From the **Manage Customer Information** window, provide the appropriate information on the **Administrator** page.

   Note: Information is required for fields with an asterisk (*).
4. Select the **System** and **Account** tabs from the **Manage Customer Information** window to provide additional information.
5. Click **OK** when you have completed the task.

Use the online Help to get additional information about customizing your account information.

Manage Event Notification

This task adds email addresses that notify you when problem events occur on your system and configures how you want to receive notification of system events from the Electronic Service Agent.

To set up notification, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Manage Event Notification**.
3. From the **Manage Event Notification** window, you can complete the following tasks:
   - Use the **Email** tab to add the email addresses that are notified when problem events occur on your system and when scheduled operations are planned for your system.
   - Use the **SNMP Trap Configuration** tab to specify locations for sending Simple Network Management Protocol (SNMP) trap messages for Hardware Management Console application programming interface events.
4. Click **OK** when you complete this task.

Use the online Help if you need additional information about this task.

Manage Connection Monitoring

Learn how to configure the timers that the connection monitoring uses to detect outages and enables or disables connection monitoring for selected machines.

You can view and, if authorized, change connection monitoring settings by machine. Connection monitoring generates serviceable events when communication problems are detected between the HMC and managed systems. If you disable connection monitoring, no serviceable events are generated for networking problems between the selected machine and this HMC.

To monitor the connections, complete the following steps:

1. In the navigation area, click the **Serviceability** icon , and then select **Service Management**.
2. In the content pane, click **Manage Connection Monitoring**.
3. From the **Manage Connection Monitoring** window, adjust the timer settings, if required, and enable or disable the server.

4. Click **OK** when you have completed the task.

Use the online Help if you need additional information about connection monitoring.

Remote operations

Connect to and use the Hardware Management Console (HMC) remotely.

Remote operations use the GUI used by a local HMC operator or the command line interface (CLI) on the HMC. You can perform operations remotely in the following ways:

- Use a remote HMC.
- Use a web browser to connect to a local HMC.
- Use an HMC remote command line.

The remote HMC is an HMC that is on a different subnet from the service processor, therefore the service processor cannot be auto discovered with IP multicast.

To determine whether to use a remote HMC or web browser that is connected to a local HMC, consider the scope of control that you need. A remote HMC defines a specific set of managed objects that are directly controlled by the remote HMC, while a web browser to a local HMC has control over the same set of managed objects as the local HMC. The communications connectivity and communications speed is an extra consideration. LAN connectivity provides acceptable communications for either a remote HMC or web browser control.

Using a remote HMC

A remote HMC gives the most complete set of functions because it is a complete HMC. Only the process of configuring the managed objects is different from a local HMC.

As a complete HMC, a remote HMC has the same setup and maintenance requirements as a local Hardware Management Console. A remote HMC needs LAN TCP/IP connectivity to each managed object (service processor) that is to be managed; therefore, any customer firewall that might exist between the remote HMC and its managed objects must allow the HMC to service processor communications to occur. A remote HMC might also need communication with another HMC for service and support. Table 10 on page 105 shows the ports that a remote HMC uses for communications.

Table 10. Ports used by a Remote HMC for Communications

Port	**Use**
udp 9900	HMC to HMC discovery
tcp 9920	HMC to HMC commands

A remote HMC needs connectivity to IBM (or another HMC that has connectivity to IBM) for service and support. The connectivity to IBM might be in the form of access to the internet (through a company firewall).

Performance and the availability of the status information and access to the control functions of the service processor depends on the reliability, availability, and responsiveness of the customer network that interconnects the remote HMC with the managed object. A remote HMC monitors the connection to each service processor and attempts to recover any lost connections and can report those connections that cannot be recovered.

Security for a remote HMC is provided by the HMC user-login procedures in the same way as a local HMC. As with a local HMC, all communication between a remote HMC and each service processor is encrypted. Certificates for secure communications are provided, and can be changed by the user if wanted.

TCP/IP access to the remote HMC is controlled through its internally managed firewall and is limited to HMC-related functions.

Using a web browser

If you need occasional monitoring and control of managed objects that are connected to a single local Hardware Management Console (HMC), use a web browser. An example of using the web browser might be an off-hours monitor from home by an operator or system programmer.

Each HMC contains a web server that can be configured to allow remote access for a specified set of users. If a customer firewall exists between the web browser and the local HMC, the ports must be accessible and the firewall setup to allow incoming requests on these ports. Table 11 on page 106 shows the ports that a web browser needs for communicating with an HMC.

Table 11. Ports that are used by a web browser for communications to the HMC

Port	Use
TCP 443	Secure (HTTPS) remote interface communication
TCP 8443	Secure browser access to web server communication
TCP 9960	Browser applet communication

[1]This port is opened in the HMC firewall when remote access is enabled in HMC Version 7.8.0 and later. This port must also be opened in any firewall that is between the remote client and the HMC.

After an HMC is configured to allow web browser access, a web browser gives an enabled user access to all the configured functions of a local HMC, except those functions that require physical access to the HMC, such as those that use the local diskette or DVD media. The user interface that is presented to the remote web browser user is the same as that of the local HMC and is subject to the same constraints as the local HMC.

The web browser can be connected to the local HMC by using a LAN TCP/IP connection and by using only encrypted (HTTPS) protocols. Logon security for a web browser is provided by the HMC user-login procedures. Certificates for secure communications are provided, and can be changed by the user.

Performance and the availability of the status information and access to the control functions of the managed objects depends on the reliability, availability, and responsiveness of the network that interconnects the web browser with the local HMC. Because there is no direct connection between the web browser and the individual managed objects, the web browser does not monitor the connection to each service processor, does not perform any recovery, and does not report any lost connections. These functions are handled by the local HMC

The web browser system does not require connectivity to IBM for service or support. Maintenance of the browser and system level is the responsibility of the customer.

If the URL of the HMC is specified by using the format https://*xxx.xxx.xxx.xxx* (where *xxx.xxx.xxx.xxx* is the IP address) and Microsoft Internet Explorer is used as the browser, a host name mismatch message is displayed. To avoid this message, a Firefox browser is used or a host name is configured for the HMC, by using the **Change Network Settings** task (see "Change Network Settings" on page 75), and this host name is specified in the URL instead of an IP address. For example, you can use the format https://*host name.domain_name* or https://*host name* (for example, by using https://hmc1.ibm.com or https://hmc1).

Preparing to use the web browser

Perform the necessary steps to prepare to use a web browser to access the Hardware Management Console (HMC).

Before you can use a web browser to access an HMC, you must complete the following tasks:

- Configure the HMC to allow remote control for specified users.
- For LAN-based connections, you must know the TCP/IP address of the HMC to be controlled, and correctly set up any firewall access between the HMC and the web browser.
- Have a valid user ID and password that is assigned by the access administrator for HMC web access.

Web browser requirements

Learn about the requirements your web browser must meet to monitor and control the Hardware Management Console (HMC).

HMC web browser support requires HTML 2.0, JavaScript 1.0, Java™ Virtual Machine (JVM), Java Runtime Environment (JRE) Version 7, and cookie support in browsers that connect to the HMC. Contact your support personnel to assist you in determining whether your browser is configured with a Java Virtual Machine. The web browser must use HTTP 1.1. If you are using a proxy server, HTTP 1.1 must be enabled for the proxy connections. Additionally, pop-up windows must be enabled for all HMCs addressed in the browser if the browser is running with pop-up windows disabled. The following browsers have been tested:

Google Chrome

HMC Version 8.1 supports Google Chrome Version 33.

Microsoft Internet Explorer

HMC Version 8.1 supports Internet Explorer 9.0, Internet Explorer 10.0, and Internet Explorer 11.0.

Note: The performance CEC task is not supported in Internet Explorer 9.0.

- If your browser is configured to use an Internet proxy, then local IP addresses are included in the exception list. For more information, see your network administrator. If you still need to use the proxy to get to the Hardware Management Console, enable Use **HTTP 1.1 through proxy connections** under the **Advanced** tab in your Internet Options window.

Mozilla Firefox

HMC Version 8.1 supports Mozilla Firefox Version 17 and Mozilla Firefox Version 24 Extended Support Release (ESR). Ensure that the JavaScript options to raise or lower windows and to move or resize existing windows are enabled. To enable these options, click the **Content** tab in the browser's Options dialog, click **Advanced** next to the Enable JavaScript option, and then select the Raise or lower windows option and the Move or resize existing windows options. Use these options to easily switch between HMC tasks. For more information about the latest Mozilla Firefox ESR levels, see Security Advisories for Firefox ESR.

Note: The following restrictions apply when you are using Mozilla Firefox while the HMC is in NIST SP 800-131a security mode:

- Mozilla Firefox cannot be used for the remote client.
- The local console cannot be used.

Other web browser considerations

Session cookies need to be enabled in order for ASMI to work when connected to HMC remotely. The ASM proxy code saves session information and uses it.

Internet Explorer

1. Click **Tools** > **Internet Options**.
2. Click the **Privacy** tab and select **Advanced**.
3. Determine whether **Always allow session cookies** is checked.
4. If not checked, select **Override automatic cookie handling** and **Always allow session cookies**.
5. For the First-party Cookies and Third-party Cookies, choose block, prompt, or accept. Prompt is preferred, in which case you are prompted every time that a site tries to write cookies. Some sites need to be allowed to write cookies.

Firefox

1. Click **Tools** > **Options**.
2. Click the **Cookies** Tab.
3. Select **Allow sites to set cookies**.
4. If you want to allow only specific sites, select **Exceptions**, and add this HMC to allow access.

Using the HMC remote command line

An alternative to performing tasks on the HMC graphical user interface is using the command line interface (CLI).

You can use the command line interface in the following situations:

- When consistent results are required. If you must administer several managed systems, you can achieve consistent results by using the command line interface. The command sequence can be stored in scripts and run remotely.
- When automated operations are required. After you develop a consistent way to manage the managed systems, you can automate the operations by starting the scripts from batch-processing applications, such as the **cron** daemon, from other systems.

On a local HMC, you can use the command line interface in the console window.

Setting up secure script execution between SSH clients and the HMC

You must ensure that your script executions between Secure Shell (SSH) clients and the Hardware Management Console (HMC) are secure.

HMCs typically are placed inside the server room where managed systems are located, so you might not have physical access to the HMC. In this case, you can remotely access it using either a remote web browser or the remote command line interface.

Note: To enable scripts to run unattended between an SSH client and an HMC, the SSH protocol must already be installed on the client's operating system.

To enable scripts to run unattended between an SSH client and an HMC, complete the following steps:

1. Enable remote command execution. For more information, see "Enable Remote Command Execution" on page 97.
2. On the client's operating system, run the SSH protocol key generator. To run the SSH protocol key generator, complete the following steps:
   a. To store the keys, create a directory that is named `$HOME/.ssh` (either RSA or DSA keys can be used).
   b. To generate public and private keys, run the following command:

      ```
      ssh-keygen -t rsa
      ```

      The following files are created in the $HOME/.ssh directory:

      ```
      private key: id_rsa
      public key: id_rsa.pub
      ```

      The write bits for both group and other are turned off. Ensure that the private key has a permission of 600."
3. On the client's operating system, use ssh and run the `mkauthkeys` command to update the HMC user's authorized_keys2 file on the HMC by using the following command:

   ```
   ssh hmcuser@hmchostname mkauthkeys -–add <the contents of $HOME/.ssh/
   id_rsa.pub>
   ```

   Note: Double quotes (") are used in commands to ensure that the remote shell can properly process the command. For example:

   ```
   ssh "mkauthkeys hscuser@somehmchost --add 'ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDa+Zc8+hn1+
   TjEXu640LqnVNB+UsixIE3c649Cgj20gaVWnFKTjcpWVahK/duCLac/zteMtVAfCx7/ae2g5RTPu7FudF2xjs4r
   +NadVXhoIqmA53a
   NjE4GILpfe5vOF25xkBdG9wxigGtJyOKeJHzgnElP7RlEeOBijJDKo5gGE12NVfBxboChm6LtKnDxLi9ahhOYtLlFehJr
   6pV/lMAEu
   Lhd6ax1hWvwrhf/
   h5Ym6J8JbLVL3EeKbCsuG9E4iN1z4HrPkT5OQLqtvC1Ajch1ravsaQqYloMTWNFzM4Qo5O3fZbLc6RuJjtJv8C5t
   4/SZUGHZxSPnQmkuii1z9hxt hscpe@vhmccloudvm179'"
   ```

To delete the key from the HMC, you can use the following command:

```
ssh hmcuser@hmchostname mkauthkeys --remove joe@somehost
```

To enable passwords that prompts for all hosts that access the HMC through SSH, use the `scp` command to copy the key file from the HMC: `scp hmcuser@hmchostname:.ssh/authorized_keys2 authorized_keys2`

Edit the `authorized_keys2` file and remove all lines in this file and then, copy it back to the HMC: `scp authorized_keys2 hmcuser@hmchostname:.ssh/authorized_keys2`

Enabling and disabling HMC remote commands

You can enable or disable the remote command line interface access to the Hardware Management Console (HMC).

To enable or disable remote commands, complete the following steps:

1. In the navigation area, select the managed system and click the **Users and Security** icon 🔓, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Command Execution**.
3. From the **Enable Remote Command Execution** window, select from the following options:
   - To enable remote commands, select **Enable remote command execution using the ssh facility**.
   - To disable remote commands, make sure **Enable remote command execution using the ssh facility** is not selected.
4. Click **OK**.

Logging in to the HMC from a LAN-connected web browser

Log in to the Hardware Management Console (HMC) remotely from a LAN-connected web browser.

Use the following steps to log in to the HMC from a LAN-connected web browser:

1. Ensure that your web browser has LAN connectivity to the wanted HMC.
2. From your web browser, enter the URL of the wanted HMC in the format `https://hostname.domain_name` (for example: `https://hmc1.ibm.com`) or `https://xxx.xxx.xxx.xxx`.

   If this connection is the first access of the HMC for the current web browser session, you might receive a certificate error. This certificate error is displayed if any of the following conditions occur:

   - The web server that is contained in the HMC is configured to use a self-signed certificate and the browser is not configured to trust the HMC as an issuer of certificates.
   - The HMC is configured to use a certificate that is signed by a certificate authority (CA) and the browser is not configured to trust this CA.

   In either case, if you know that the certificate that is being displayed to the browser is the one used by the HMC, you can continue and all communications to the HMC is encrypted.

   If you do not want to receive notification of a certificate error for the first access of any browser session, you can configure the browser to trust the HMC or the CA. In general, to configure the browser, use one of the following methods:

   - You must indicate that the browser permanently trusts the issuer of the certificate.
   - By viewing the certificate and installing to the database of trusted CAs, the certificate of the CA that issues the certificate is used by the HMC.

   If the certificate is self-signed, the HMC itself is considered the CA that issues the certificate.
3. When prompted, enter the user name and password that is assigned by your administrator.

Managing OpenBMC-based and BMC-based systems by using the HMC

Learn how to manage OpenBMC-based and BMC-based systems by using the Hardware Management Console (HMC).

About this task

Learn about the tasks that you perform from the console and how to navigate the baseboard management controller (BMC) by using the web-based user interface with graphical views of managed systems and simplified navigation.

Note: You cannot manage OpenBMC-based and BMC-based systems while the HMC is running in NIST mode.

Add Managed Systems

Learn how to add a managed Baseboard Management Controller (BMC) system to the Hardware Management Console (HMC).

To add one or more managed BMC systems to the HMC, complete the following steps:

1. From the HMC dashboard, click **Connect Systems**
2. From the **Add Managed Systems** window, you can add a BMC system by completing the following fields:
   - **IP Address/Host name**
   - **Username (BMC system)**

     Notes:

     – For model 8335-GTH and 8335-GTX, the default user name is `admin`.

     – For model 9006-12P and 9006-22P, the default user name is `ADMIN`.
   - **Password**

   Alternatively, you can specify a range of IP addresses and click **OK** to view a list of systems that were discovered. You can select one or more discovered systems to add to the HMC.

   Note: The discovery process can take a long time to complete.
3. Click **OK** to add the managed system to the HMC.

Use the online Help if you need additional information about this task.

Systems Management for Servers

Systems Management displays tasks to manage servers. Use these tasks to set up, configure, view status, troubleshoot, and apply solutions for servers.

These tasks are listed when a managed system is selected. The tasks that are listed in the menu pod change as selections are made in the work area.

Operations

Operations contains the tasks for operating managed systems.

Power Off

Shut down the managed system.

Choose from the following options:

Normal power off

The Normal power off mode shuts down the system's operations in a controlled manner. During the shutdown, programs running active jobs are allowed to perform cleanup (end-of-job processing).

Power On

Use the **Power On** task to start a managed system.

Choose from the following options to power on your managed system:

Normal: Select this option to specify that the HMC uses the current setting for the partition start policy to determine how to power on the managed system. The default setting is set to the following value:

- **Auto-Start Always**: This option specifies that the HMC power on logical partitions automatically after the managed system powers on. If powering on the managed system is the result of a user action, the HMC starts all partitions that are configured for automatic startup. If powering on the managed system is the result of an automatic recovery process, the HMC starts only those logical partitions that were running at the time the system is powered off. This option is always available for selection.

Schedule Operations

Create a schedule for certain operations to be performed on the managed system without operator assistance.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule power on or off operations for a managed system.

The Scheduled Operations task displays the following information for each operation:

- The processor that is the object of the operation.
- The scheduled date
- The scheduled time
- The operation
- The number of remaining repetitions

From the **Scheduled Operations** window, you can perform the following tasks:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

You can schedule an operation to occur once or you can schedule it to repeat. You must provide the time and date that you want the operation to occur. If you want the operation to repeat, you are asked to select the following options:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)
- The total number of repetitions. (required)

The operations that you can schedule for the managed system include the following operations:

Power Off Managed System

Schedules an operation for a system power off at regular intervals for a managed system.

Power On Managed System

Schedules an operation for a system power-on at regular intervals for a managed system.

To schedule operations on the managed system, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Servers**.
2. In the content pane, select one or more managed systems.
3. In the menu pod, select **Actions** > **View All Actions** > **Operations** > **Schedule Operations**.
4. From the **Scheduled Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, click **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
5. To return to the HMC workplace, point to **Operations** and then click **Exit**.

Launch BMC System Management

The Hardware Management Console (HMC) can connect directly to the Baseboard Management Controller (BMC) for a selected system.

The BMC system management is an interface to the service processor that allows you to manage the operation of the server, such as auto power restart, and to view information about the server, such as the error log and vital product data.

To connect to the BMC, complete the following steps:

Note: To access the BMC user interface, you must be at the console or have access to the BMC by using a supported web browser.

1. In the navigation area, click the **Resources** icon , and then select **All Servers**.
2. In the content pane, select one or more managed systems.
3. In the menu pod, select **Actions** > **View All actions** > **Operations** > **Launch BMC System Management**.
4. Click **Continue**.

Configuring Call Home

Problems on your BMC-based managed system are reported to the Hardware Management Console (HMC) as events. You can set up alerts to be automatically notified of any events.

Note: You must enable SNMP traps in the HMC to receive alerts. To enable SNMP traps, navigate to **Console Settings** > **Change Network Settings** > **LAN Adapters** > **Details** > **Firewall Settings**. Select **SNMP Traps** and **SNMP Agent** from the table and then click **Allow Incoming**.

To set up alerts for call home, complete the following steps:

Note: This procedure is applicable for model 9006-12P, 9006-22C, and 9006-22P.

1. From the **BMC System Management** window, click **Configuration** > **Alerts**.
2. Select any alert from the table and click **Modify**.

Note: You can set up multiple HMCs to receive traps. Duplicate reporting of events by multiple HMCs is possible as duplicate event verification is not performed.

3. Complete the following fields:
   - **Event Severity**
   - **Destination IP**
4. Click **Save**.
5. Verify the new alert in the table.

Use the online Help if you need additional information about this task.

Rebuild

You can extract the configuration information from the managed system and rebuild the information on the Hardware Management Console (HMC).

This task does not disrupt the operation of the running server.

Rebuilding the managed system updates the information on the HMC about the managed system. Rebuilding the managed system is useful when the state of the managed system is `Incomplete`. The Incomplete state means that the HMC cannot gather complete information from the managed system about logical partitions, profiles, or resources.

Rebuilding the managed system is different from refreshing the **HMC** window. When the managed system is rebuilt, the HMC extracts the information from the managed system. You cannot start other tasks while the HMC rebuilds the managed system. This process can take several minutes.

Updates

Display tasks to view system information, manage updates on your Hardware Management Console (HMC), or check system readiness.

Change Licensed Internal Code

Change the Licensed Internal Code of a managed BMC system by using your Hardware Management Console (HMC).

The system firmware is a combination of the BMC firmware and the PNOR firmware. You must update both the BMC firmware and the PNOR firmware for the system to operate properly. If you update only one type of firmware, but do not update the other type of firmware, system errors might occur.

To change the Licensed Internal Code, complete the following steps:

1. In the navigation area, click the **Resources** icon , and then select **All Servers**.
2. Select the server for which you want to view system information.
3. In the menu pod, expand **Actions** and then expand **Updates**.
4. Select **Change Licensed Internal Code** > **BMC Change Licensed Internal Code**.
5. Follow the onscreen instructions in the **BMC Change Licensed Internal Code** guided wizard.

   Note: The BMC system must be in the powered off state before you can proceed with the wizard.
6. When you complete this task, click **Close**.

Use the online Help if you need additional information about this task.

Attention LED

View system attention LED information, light specific LEDs to identify a system component, and test all LEDs on a managed system.

The system provides several LEDs that help identify various components, such as enclosures or field replaceable units (FRUs), in the system. For this reason, they are called **Identify** LEDs. Individual LEDs are on or near the components. The LEDs are located either on the component itself or on the carrier of the component (for example, memory card, fan, memory module, or processor). LEDs are either green or amber. Green LEDs indicate either of the following states:

- Electrical power is present.
- Activity is occurring on a link. (The system might be sending or receiving information).

Amber LEDs indicate a fault or identify condition. If your system or one of the components on your system has an amber LED turned on or flashing, identify the problem and take the appropriate action to restore the system back to normal.

You can activate or deactivate the following types of identify LEDs:

Identify LED for an enclosure

If you want to add an adapter to a specific drawer (enclosure), you need to know the machine type, model, and serial number (MTMS) of the drawer. To determine whether you have the correct MTMS for the drawer that needs the new adapter, you can activate the LED for a drawer and verify that the MTMS corresponds to the drawer that requires the new adapter.

You can deactivate a system attention LED. For example, you might determine that a problem is not a high priority and decide to repair the problem later. However, you want to be alerted if another problem occurs, so you must deactivate the system attention LED so that it can be activated again if another problem occurs.

Connections

You can view the Hardware Management Console (HMC) connection status to service processors, reset those connections, connect another HMC to the selected managed system, or disconnect another HMC.

If you select a managed system in the work area, the following tasks pertain to that managed system.

Service Processor Status

View information about the status of the Hardware Management Console (HMC) connection to the service processors on the managed system.

About this task

To show the service processor connection status to the service processors on the managed system, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Servers**.
2. Select the server for which you want to view service processor connection status.
3. In the menu pod, select **Actions** > **View All Actions** > **Connections** > **Service Processor Status**.

Reset or Remove Connections

Reset or remove a managed system from the Hardware Management Console (HMC) interface.

About this task

To reset or remove connections, complete the following steps:

Procedure

1. In the navigation area, click the **Resources** icon , and then select **All Servers**.
2. Select the server that you want to reset or remove.

3. In the menu pod, select **Actions** > **View All Actions** > **Connections** > **Reset or Remove Connections**.
4. Select **Reset Connection** or **Remove Connection**.
5. Click **OK**.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s user name and IP address for purposes of session management. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled “Cookies, Web Beacons and Other Technologies”.

Programming interface information

This Managing the Hardware Management Console publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM Hardware Management Console Version 9 Release 2 Maintenance Level 950.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM®

Power Systems

Problem analysis, system parts, and locations for the IBM Power Systems HMC (7063-CR2)

IBM

Note

Before using this information and the product it supports, read the information in “Safety notices” on page v, “Notices” on page 25, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Power Systems servers that contain the POWER9™ processor and to all associated models.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.

- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- 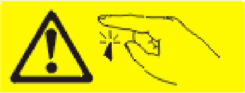 Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).

- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.

  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer’s DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

⚠ **CAUTION:**

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.

  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

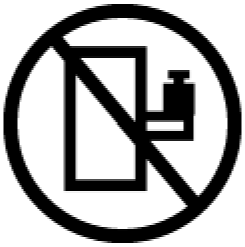

⚠ **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.

- Before extending the rack to the installation position, read the installation instructions.
- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

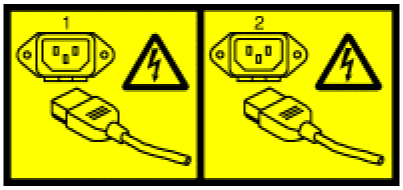

or

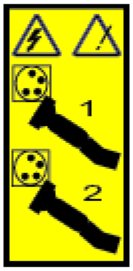

or

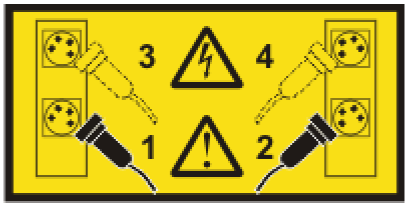

or

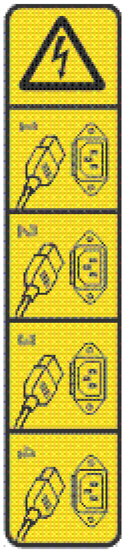

or

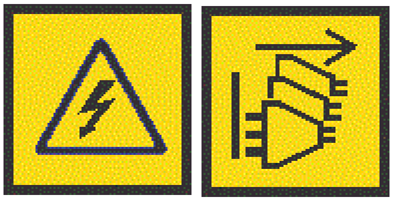

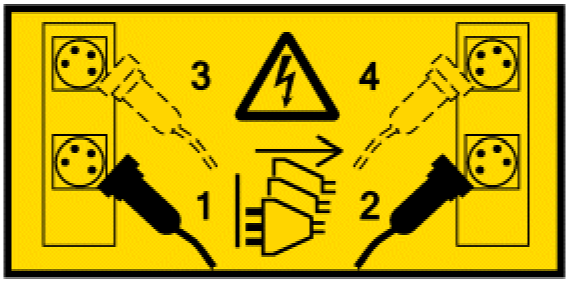

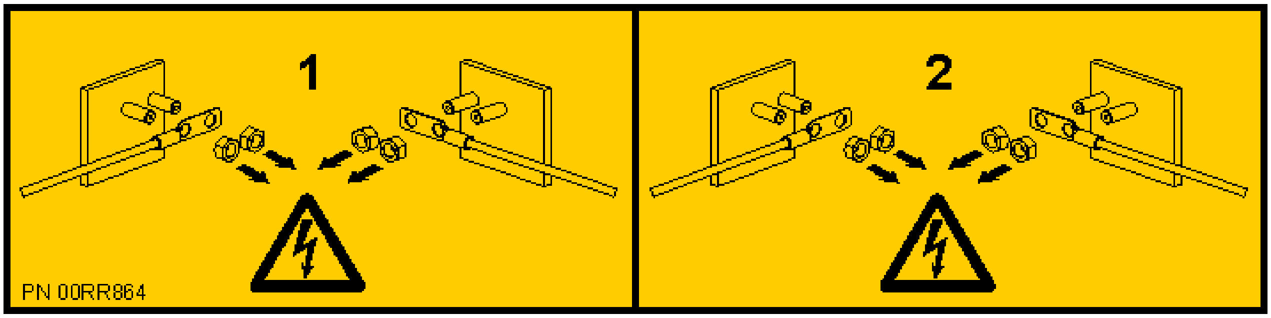

DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

CAUTION: A hot surface nearby. (L007)

(L008)

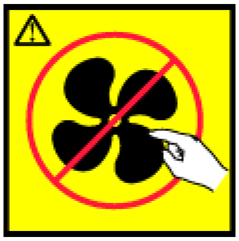

CAUTION: Hazardous moving parts nearby. (L008)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

CAUTION: This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

CAUTION: This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

CAUTION: Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

CAUTION: The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)

CAUTION: Regarding IBM provided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.
- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four

(4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.

- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.
- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intrabuilding ports of this equipment are suitable for connection to intrabuilding or unexposed wiring or cabling only. The intrabuilding ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The ac-powered system does not require the use of an external surge protection device (SPD).

The dc-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The dc-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Beginning troubleshooting and problem analysis

This information provides a starting point for analyzing problems.

This information is the starting point for diagnosing and repairing systems. From this point, you are guided to the appropriate information to help you diagnose problems, determine the appropriate repair action, and then complete the necessary steps to repair the system.

Notes:

- Update the system firmware to the latest level before you start problem analysis. If you update the system firmware, you have the latest available fixes and improvements for error handling, reporting, and isolation. For instructions about updating the system firmware, see Getting fixes.
- Some service procedures use OpenBMC tool commands. To download and install the OpenBMC tool, see Downloading and installing the OpenBMC tool.

What type of problem are you dealing with?	Problem analysis procedure
You do not know the type of problem.	Go to "Determining the problem analysis procedure to perform" on page 1.
You have an FQPSP*xxxxxxx* event code.	Go to FQPSP*xxxxxxx* Event Codes.
A baseboard management controller (BMC) access problem occurred.	Go to "Resolving a BMC access problem" on page 2.
The system does not power on (the power button or the BMC power on command does not power on the system).	Go to "Resolving a power problem" on page 4.
A system firmware boot failure occurred (the system started but was not able to boot to the Petitboot menu).	Go to "Resolving a system firmware boot failure" on page 5.
A video graphics array (VGA) monitor problem occurred (the system started but no video is displayed on the monitor).	Go to "Resolving a VGA monitor problem" on page 6.
An operating system boot failure occurred (the system booted to the Petitboot menu but the operating system did not start).	Go to "Resolving an operating system boot failure" on page 6.
A processor, memory, power, or cooling hardware failure occurred.	Go to "Resolving a hardware problem" on page 7.
Missing or faulty PCIe adapter or device.	Go to "Resolving a PCIe adapter or device problem" on page 7.

Determining the problem analysis procedure to perform

Learn how to identify the correct problem analysis procedure to perform.

About this task

To determine the correct problem analysis procedure to perform, complete the following steps:

Procedure

1. After you apply power to the system, are the power supply LEDs green?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a power problem" on page 4.

2. Can you access the baseboard management controller (BMC) across the network?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a BMC access problem" on page 2.

3. Can you boot the system to the Petitboot menu?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a system firmware boot failure" on page 5.

4. Is video displayed on the video graphics array (VGA) monitor?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a VGA monitor problem" on page 6.

5. Can you start the operating system?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving an operating system boot failure" on page 6.

6. Go to "Resolving a hardware problem" on page 7. **This ends the procedure.**

Resolving a BMC access problem

Learn how to identify the service action that is needed to resolve a baseboard management controller (BMC) access problem.

Procedure

1. Ensure that the BMC password is not set to the default password. For information about changing the default password, see Logging on to the OpenBMC GUI. Does the problem persist?

If	Then
Yes:	Continue with the next step.
No:	**This ends the procedure.**

2. Are both ends of the network cable seated securely?

If	Then
Yes:	Continue with the next step.
No:	Seat both ends of the cable securely. If the problem persists, continue with the next step.

3. Is the operating system available?

If	Then
Yes:	Continue with the next step.
No:	Continue with step "5" on page 3.

4. Verify that the BMC network settings are correct.

   a) In the navigation area, click the **HMC Management** icon , and then select **Console Settings**.

   b) In the content pane, click **Change BMC/IPMI network settings**.

   c) Verify that the MAC address and the IP address settings are correct.

   Does the BMC access problem persist?

If	Then
Yes:	Continue with the next step.
No:	**This ends the procedure.**

5. Power off the system and disconnect all AC power cords for 30 seconds. Then, reconnect the AC power cords and power on the system. Does the BMC access problem persist?

If	Then
Yes:	Continue with the next step.
No:	**This ends the procedure.**

6. Verify that the BMC network settings are correct.

   Note: To verify the BMC network settings, you must have a cabled serial connection or a monitor and keyboard.

   a) Power on the system by using the power button on the front of the system. Wait 1 - 2 minutes for the system to display the Petitboot menu.

   b) When the Petitboot menu is displayed, press any key to interrupt the boot process. Then, select Exit to Shell.

   c) For a shared BMC Ethernet port, type the following command and press Enter:

```
ipmitool lan print 1
```

   For a dedicated BMC Ethernet port, type the following command and press Enter:

```
ipmitool lan print 2
```

   To determine the location of the shared and dedicated BMC Ethernet ports, see Table 1 on page 3.

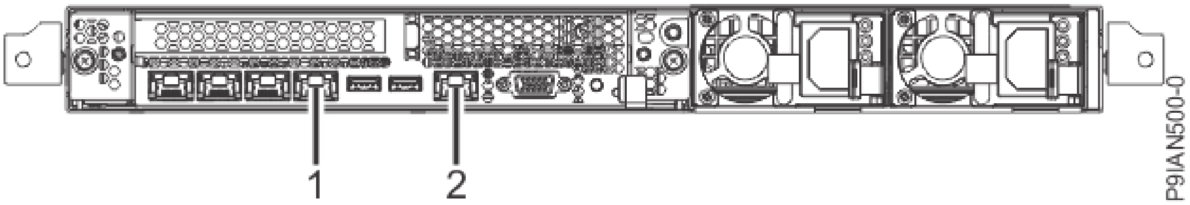

Figure 1. Rear BMC Ethernet ports

Table 1. BMC Ethernet ports

Identifier	Description
1	Shared BMC Ethernet
2	Dedicated BMC Ethernet

   d) Verify that the MAC address and the IP address settings are correct. Then, continue with the next step.

Note: If the IP address setting is incorrect, go to Configuring the BMC IP address. If the MAC address is 00:00:00:00:00:00, go to “Contacting IBM service and support” on page 12.

7. Complete the following actions:

   a. Power on to the Petitboot menu.

   b. Update the system firmware. For instructions, see Getting fixes.

   Are you able to access the BMC?

If	Then
Yes:	**This ends the procedure.**
No:	Continue with the next step.

8. Replace the system backplane. Go to “7063-CR2 locations” on page 15 to identify the physical location and the removal and replacement procedure. **This ends the procedure.**

Resolving a power problem

Learn how to identify the service action that is needed to resolve a power problem.

Procedure

1. Is the amber LED (bottom LED) of a power supply on solid and is the amber LED on the front of the system turned off?

If	Then
Yes:	Ensure that the power cords for both power supplies are fully seated and that the power distribution units (PDUs) and power outlets are supplying electricity. **This ends the procedure.**
No:	Continue with the next step.

2. Are the power supply LEDs turned off?

If	Then
Yes:	Continue with the next step.
No:	Continue with step “4” on page 4.

3. Perform the following actions, one at a time until the problem is resolved:

   a. Ensure that all of the power cords are fully seated in the power supplies.

   b. Ensure that all of the power cords are fully seated in the power distribution units (PDUs) or wall outlets.

   c. If the power cords are plugged into PDUs, ensure that the PDUs are turned on.

   d. Ensure that all of the power cords are plugged into PDUs or wall outlets that are supplying electricity.

   e. Replace the power cords.

   f. Replace the power supplies. Go to “7063-CR2 locations” on page 15 to identify the physical location and the removal and replacement procedure. **This ends the procedure**.

4. Is the amber LED of a power supply on solid and is the amber system attention LED on the front of the system on solid?

If	Then
Yes:	Continue with the next step.
No:	Go to “Contacting IBM service and support” on page 12. **This ends the procedure.**

5. Perform the following actions, one at a time until the problem is resolved:

a. Resolve any serviceable alerts that are in the event log. Go to “Resolving a hardware problem” on page 7.

b. Ensure that the power supply is fully seated in the system.

c. Ensure that the power supply fan is not blocked.

d. Replace the power supply. Go to “7063-CR2 locations” on page 15 to identify the physical location and the removal and replacement procedure. **This ends the procedure**.

Resolving a system firmware boot failure

Learn how to identify the service action that is needed to resolve a failure while booting your system firmware.

Procedure

1. After you press the power button, did the system turn on but fail to display the Petitboot menu?

If	**Then**
Yes:	Continue with the next step.
No:	Continue with step “6” on page 5.

2. Does the baseboard management controller (BMC) respond to commands?

   Note: To determine whether the BMC responds to commands, run the following OpenBMC tool command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis power status
```

If	**Then**
Yes:	Continue with the next step.
No:	Continue with step “4” on page 5.

3. Complete the following actions:

   a. Update the system firmware. For instructions, see Getting fixes.

   b. Check the system event logs. For instructions, see “Identifying a service action by using system event logs” on page 10. Then, continue with step “6” on page 5.

4. Disconnect the power cords from the system for 30 seconds. Reconnect the power cords, wait 5 minutes, and then continue with the next step.

5. Does the baseboard management controller (BMC) respond to commands?

   Note: To determine whether the BMC responds to commands, run the following OpenBMC tool command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis power status
```

If	**Then**
Yes:	Update the system firmware. For instructions, see Getting fixes. **This ends the procedure.**
No:	Replace the system backplane. Go to “7063-CR2 locations” on page 15 to identify the physical location and the removal and replacement procedure. **This ends the procedure.**

6. Power off the system and disconnect all AC power cords for 30 seconds. Then, reconnect the AC power cords and power on the system. Does the system boot successfully?

If	Then
Yes:	**This ends the procedure.**
No:	Go to "Resolving a hardware problem" on page 7. **This ends the procedure.**

Resolving a VGA monitor problem

Learn how to identify the service action that is needed to resolve a video graphics array (VGA) monitor problem.

Procedure

1. Is the system powered on and is the VGA monitor connected to the VGA display port, but no video is displayed?

If	Then
Yes:	Continue with the next step.
No:	Power on the system. **This ends the procedure.**

2. Complete the following steps, one at a time until the problem is resolved:
   a) Ensure that the network image that is specified is a valid boot image.
   b) Ensure that the VGA cable is properly seated to the server port and to the monitor port.
   c) Verify that your monitor and your VGA cable are working properly by testing them on a system that is known to be working properly. If the monitor or the VGA cable does not work properly, replace it.
   d) Verify that the system is powered on by activating a serial over LAN (SOL) session through the baseboard management controller (BMC). If the system is not active, go to "Resolving a system firmware boot failure" on page 5.
   e) Replace the system backplane. Go to "7063-CR2 locations" on page 15 to identify the physical location and the removal and replacement procedure. **This ends the procedure**.

Resolving an operating system boot failure

Learn how to identify the service action that is needed to resolve a failure while booting your operating system.

Procedure

1. Was the system recently installed, serviced, moved, or upgraded?

If	Then
Yes:	Ensure that all cables are properly seated in the connection path to the designated boot device. **This ends the procedure.**
No:	Continue with the next step.

2. Petitboot displays all recognized bootable images to use by default. Is the boot image recognized by Petitboot?

If	Then
Yes:	Continue with the next step.
No:	Select the Petitboot menu option to refresh the boot images. If the problem persists, go to "Resolving a storage device problem" on page 10. **This ends the procedure.**

3. Does an operating system error occur during the boot?

If	Then
Yes:	Recover the operating system with the tools provided for the operating system. If that does not resolve the problem, reinstall the operating system. **This ends the procedure.**
No:	Reinstall the operating system. **This ends the procedure.**

Resolving a hardware problem

Learn how to identify the service action that is needed to resolve a hardware problem.

Procedure

1. If you have not already done so, manually boot the system.
2. Go to "Identifying a service action by using system event logs" on page 10. Then, continue with the next step.
3. Was a service action identified?

If	Then
Yes:	Continue with the next step.
No:	Go to step "5" on page 7.

4. Did the service action fix the problem?

If	Then
Yes:	**This ends the procedure.**
No:	Go to step "5" on page 7.

5. Go to "Resolving a PCIe adapter or device problem" on page 7. Then, continue with the next step.
6. Was a service action identified?

If	Then
Yes:	Continue with the next step.
No:	Go to "Collecting diagnostic data" on page 11. Then, go to "Contacting IBM service and support" on page 12. **This ends the procedure.**

7. Did the service action fix the problem?

If	Then
Yes:	**This ends the procedure.**
No:	Go to "Collecting diagnostic data" on page 11. Then, go to "Contacting IBM service and support" on page 12. **This ends the procedure.**

Resolving a PCIe adapter or device problem

Learn how to access log files, information to identify types of events, and a list of potential problems and service actions.

Procedure

1. To identify the correct service procedure to perform by using operating system log information, complete the following steps:

   a) Log in as the **hscroot** user.

   b) To display the operating system logs, type `less /var/log/messages` and press **Enter**.

2. Scan the operating system logs that occurred around the time that the problem started for the first occurrence of keywords, such as fail, failure, or failed. When you find a keyword that accompanies one or more of the resource names in the following table, a service action is required. Use the following table to determine the service procedure to perform for your type of problem.

Table 2. Resource names, examples, and service procedures for different types of operating system logs.

Resource name	Example of a log requiring a service action	Type of problem	Service procedure
eth1, eth2, eth3, enP*xxxxx*, where *xxxxx* indicates the network port.	`Failed to re-initialize device`	Network	Go to "Resolving a network adapter problem" on page 9.

tg3	`PCI I/O error detected.`
	`Link is Down`
		Network	Go to "Resolving a network adapter problem" on page 9.

sda, sdb, sdc	`FAILED Result`	Storage	Go to "Resolving a storage device problem" on page 10.
EEH	`Detected error on PHB#`*xxx*, where *xxx* is the PHB number.	PCIe bus or adapter	Resolve any device driver errors that are related to I/O and that occurred near the time of this operating system log entry.
	xxx `has failed 6 times in the last hour and has been permanently disabled`, where *xxx* is the PCI bus number.	PCIe bus or adapter	Ensure that the correct device drivers are properly installed for the device. If the problem persists, replace the adapter in the PCIe slot that is specified in the operating system log entry.

Resolving a network adapter problem

Learn about the possible problems and service actions that you can perform to resolve a network adapter problem.

About this task

Table 3. Network adapter problems and service actions

Problem	Service action
System unable to find adapter	1. Verify that the most recent firmware is installed on the system. Otherwise, install the most recent firmware if it is not already installed. 2. Restart the system. 3. Replace the adapter. 4. Replace the system backplane. 5. Replace the central processing unit (CPU).
Adapter stops working suddenly	1. If the system was recently installed, moved, serviced, or upgraded, verify that the adapter is seated properly and all associated cables are correctly connected. 2. Inspect the PCIe socket and verify that there is no dirt or debris in the socket. 3. Inspect the card and verify that it is not physically damaged. 4. Verify that all cables are properly seated and are not physically damaged. 5. Replace the adapter. 6. Replace the system backplane. 7. Replace the CPU.
Link indicator light on the adapter is off	1. Verify that the cable functions properly by testing it with a known working connection. 2. Verify that the port or ports on the switch are enabled and functional. 3. Verify that the switch and adapter are compatible. 4. Replace the adapter.
Link light on the adapter is on, but there is no communication from the adapter	1. Verify that the most recent driver is installed, or install the most recent driver if it is not already installed. 2. Verify that the adapter and its link have compatible settings, such as speed and duplex configuration.

Resolving a storage device problem

Learn about the possible problems and service actions that you can perform to resolve a storage device problem.

About this task

Note: To determine the location of the storage device, see Removing and replacing a drive in the 7063-CR2.

Table 4. Storage device problems and service actions

Problem	**Service action**
System unable to find a storage device that is at the front of the system	1. If the system was recently installed, moved, serviced, or upgraded, verify that the device is seated and installed properly. 2. Verify that the device is compatible with your system. 3. Verify that all internal cables are properly seated and are not physically damaged. 4. Verify that the most recent firmware is installed on the system. Otherwise, install the most recent firmware if it is not already installed. 5. Replace the drive. 6. Replace the cable. 7. Replace the drive holder.
Drive stops working suddenly	1. Verify that all internal cables are properly seated and are not physically damaged. 2. Check the system logs to verify whether the system detected a problem. 3. Replace the drive. 4. Replace the cable.
Other problems	Check the messages and resolve any other problems that were detected. Then, test the drive again. If the drive continues not to function, refer to the documentation for the drive.

Identifying a service action by using system event logs

Use the OpenBMC tool to examine system event logs (SELs) to identify a service action.

Procedure

1. From a system that has the OpenBMC tool installed, type the following command and press Enter:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel print
```

2. Is there an **Active Alerts** section displayed in the output of the command?

If	**Then**
Yes:	Continue with the next step.
No	No service action is required. **This ends the procedure.**

If	Then
	Note: Alerts that are displayed in the **Historical Alerts** section do not require service.

3. Is there an entry in the **Active Alerts** section with a value of **Yes** in the **Serviceable** column?

If	Then
Yes:	Continue with the next step.
No	No service action is required. **This ends the procedure.** **Note:** Alerts with a value of **No** that are displayed in the **Serviceable** column do not require service.

4. Starting with the first entry in the **Active Alerts** section with a value of **Yes** in the **Serviceable** column, complete the following steps until all entries are resolved:
   a. Record the log number that is displayed in the **Entry** column.
   b. Record the FQPSP*xxxxxxx* value that is displayed in the **ID** column. Then, go to FQPSP*xxxxxxx* Event Codes and complete the service action that is indicated for the FQPSP*xxxxxxx* event code.
   c. After the service action is complete and the problem is resolved, type the following command and press Enter:

   `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel resolve -n x`, where *x* is the log number that you recorded in step "4.a" on page 11.

This ends the procedure.

Verifying a repair

Learn how to verify hardware operation after you make repairs to the system.

Procedure

1. Power on the system.
2. Scan the system event logs (SELs) for serviceable events that occurred after system hardware was replaced. For information about SELs that require a service action, see "Identifying a service action by using system event logs" on page 10.
3. Did any serviceable SEL events occur after hardware was replaced?

If	Then
Yes:	The problem is not resolved. Go to "Identifying a service action by using system event logs" on page 10 and complete the service actions indicated. **This ends the procedure.**
No:	The problem is resolved. **This ends the procedure.**

Collecting diagnostic data

Learn how to collect diagnostic data to send to IBM service and support.

About this task

To collect diagnostic data, complete the following steps:

Procedure

1. Are you able to log on to the Hardware Management Console (HMC)?

If	Then
Yes:	Continue with the next step.
No:	Go to step "3" on page 12.

2. Collect diagnostic data from the 7063-CR2 by using the **PEDBG** command on the HMC. To collect diagnostic data, go to HMC Enhanced View: Collecting PEDBG from the HMC (http://www.ibm.com/support/pages/hmc-enhanced-view-collecting-pedbg-hmc) and complete the steps that are indicated. Send the data that you collected during this procedure to IBM service and support. **This ends the procedure.**
3. Can you boot the system to the Petitboot menu or is another system available that has the Linux® operating system?

If	Then
Yes:	Continue with the next step.
No:	Go to "Contacting IBM service and support" on page 12.

4. To collect system event logs, complete the following steps:
   a) Go to the IBM Support Portal (http://www.ibm.com/mysupport/s/).
   b) In the search field, type `Scale-out LC System Event Log Collection Tool`.
   c) Click the **Scale-out LC System Event Log Collection Tool** entry and follow the instructions to install and run the system event log collection tool. Then, continue with the next step.
5. Send the data that you collected during this procedure to IBM service and support. **This ends the procedure.**

Contacting IBM service and support

You can contact IBM service and support by telephone or through the IBM Support Portal.

Before you contact IBM service and support, go to "Beginning troubleshooting and problem analysis" on page 1 and complete all of the service actions indicated. If the service actions do not resolve the problem, or if you are directed to contact support, go to "Collecting diagnostic data" on page 11. Then, use the information below to contact IBM service and support.

Customers in the United States, United States territories, or Canada can place a hardware service request online. To place a hardware service request online, go to the IBM Support Portal (http://www.ibm.com/support/entry/portal/product/power/scale-out_lc).

For up-to-date telephone contact information, go to the Directory of worldwide contacts website (www.ibm.com/planetwide/).

Table 5. Service and support contacts

Type of problem	Call
• Advice • Migrating • "How to" • Operating • Configuring • Ordering • Performance • General information	• 1-800-IBM-CALL (1–800–426–2255) • 1-800-IBM-4YOU (1–800–426–4968)

Table 5. Service and support contacts (continued)

Type of problem	Call
Software: • Fix information • Operating system problem • IBM application program • Loop, hang, or message Hardware: • IBM system hardware broken • Hardware reference code • IBM input/output (I/O) problem • Upgrade	1-800-IBM-SERV (1–800–426–7378)

Finding parts and locations

Locate physical part locations and identify parts with system diagrams.

Locate the FRU

Use the graphics and tables to locate the field-replaceable unit (FRU) and identify the FRU part number.

7063-CR2 locations

Use this information to find the location of a field-replaceable unit (FRU) in the system unit.

Rack views

The following diagrams show FRU layouts in the system. Use these diagrams with the following tables.

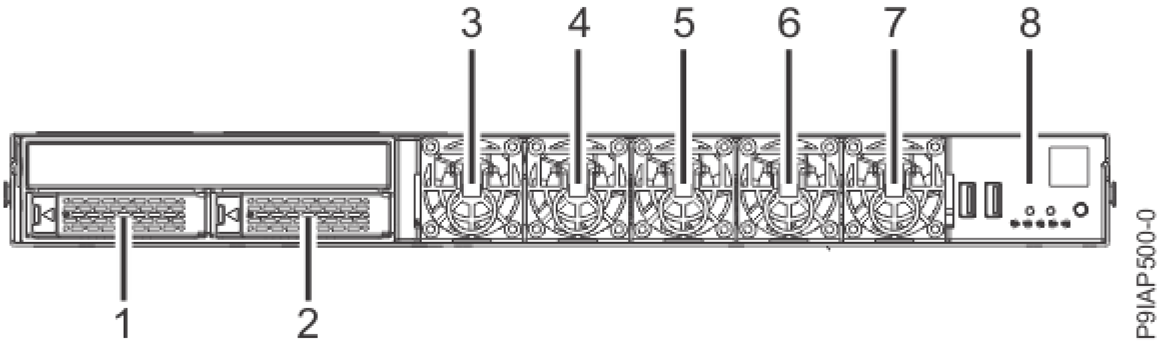

Figure 2. Front view

Table 6. Front view locations

Index number	FRU description	FRU removal and replacement procedures
1	Drive 0	See Removing and replacing a drive in the 7063-CR2.
2	Drive 1	
3	Fan 0	See Removing and replacing fans in the 7063-CR2.
4	Fan 1	
5	Fan 2	
6	Fan 3	
7	Fan 4	
8	Control panel	See Removing and replacing the control panel in the 7063-CR2.

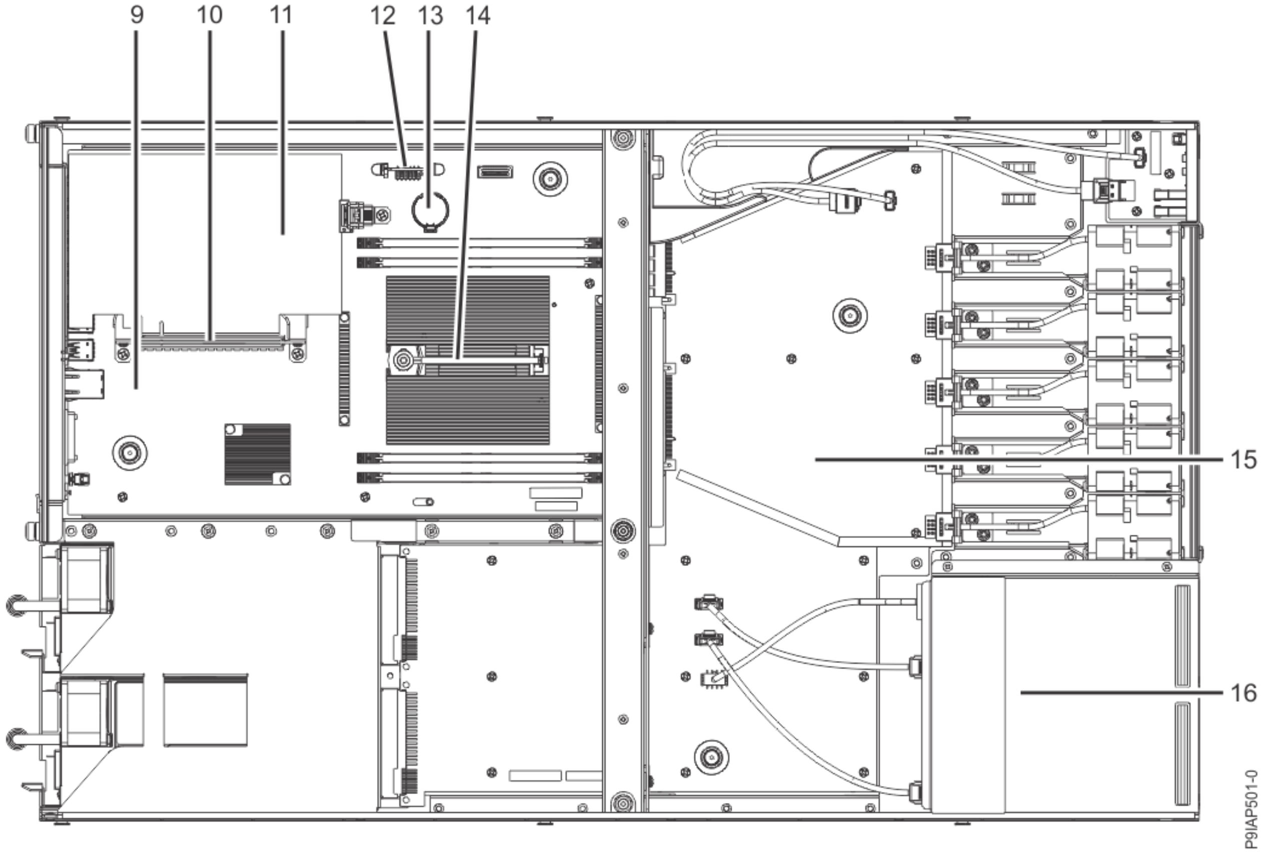

Figure 3. Top view

Table 7. Top view locations

Index number	FRU description	FRU removal and replacement procedures
9	System backplane	See Removing and replacing the system backplane in the 7063-CR2.
10	Riser	See Removing and replacing the PCIe riser in the 7063-CR2.
11	PCIe adapter	See Removing and replacing PCIe adapters in the 7063-CR2.
12	Trusted platform module	See Removing and replacing the trusted platform module in the 7063-CR2.
13	Time-of-day battery	See Removing and replacing the time-of-day battery in the 7063-CR2.
14	CPU 0	See Removing and replacing the system processor module in the 7063-CR2.
15	Power distribution board	See Removing and replacing the power distribution board in the 7063-CR2.

Table 7. Top view locations (continued)		
Index number	**FRU description**	**FRU removal and replacement procedures**
16	Drive holder	See Removing and replacing the drive holder in the 7063-CR2.

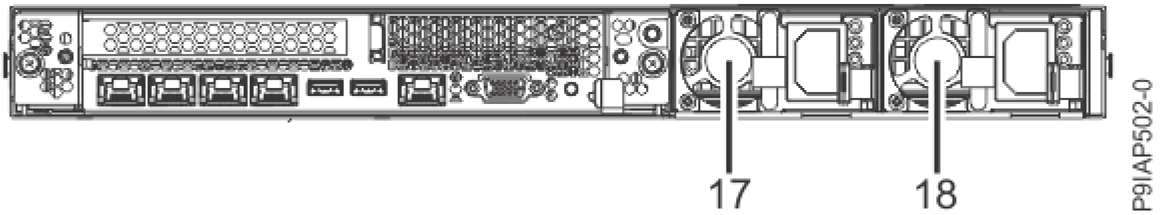

Figure 4. Rear view

Table 8. Rear view locations		
Index number	**FRU description**	**FRU removal and replacement procedures**
17	Power supply unit 1 (PSU 1)*	See Removing and replacing a power supply in the 7063-CR2.
18	Power supply unit 0 (PSU 0)*	

* The E*x* labels on the chassis do not match the power supply unit number.

Memory locations

The following diagram shows memory DIMMs and their corresponding field-replaceable unit (FRU) layouts in the system. Use this diagram with the following table.

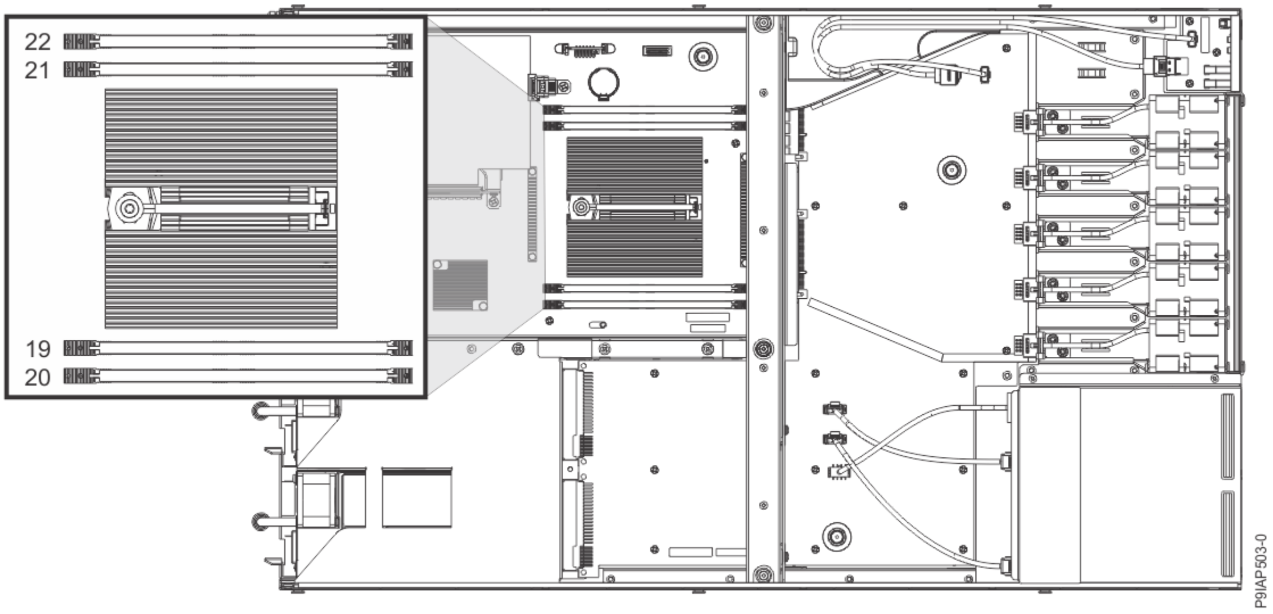

Figure 5. Memory locations

Table 9. Memory locations

Index number	FRU description	FRU removal and replacement procedures
19	DIMM 0	See Removing and replacing memory in the 7063-CR2.
20	DIMM 1	
21	DIMM 2	
22	DIMM 3	

7063-CR2 parts

Use this information to find the field-replaceable unit (FRU) part number.

After you identify the part number of the part that you want to order, go to Advanced Part Exchange Warranty Service. Registration is required. If you are not able to identify the part number, go to Contacting IBM service and support.

Rack final assembly

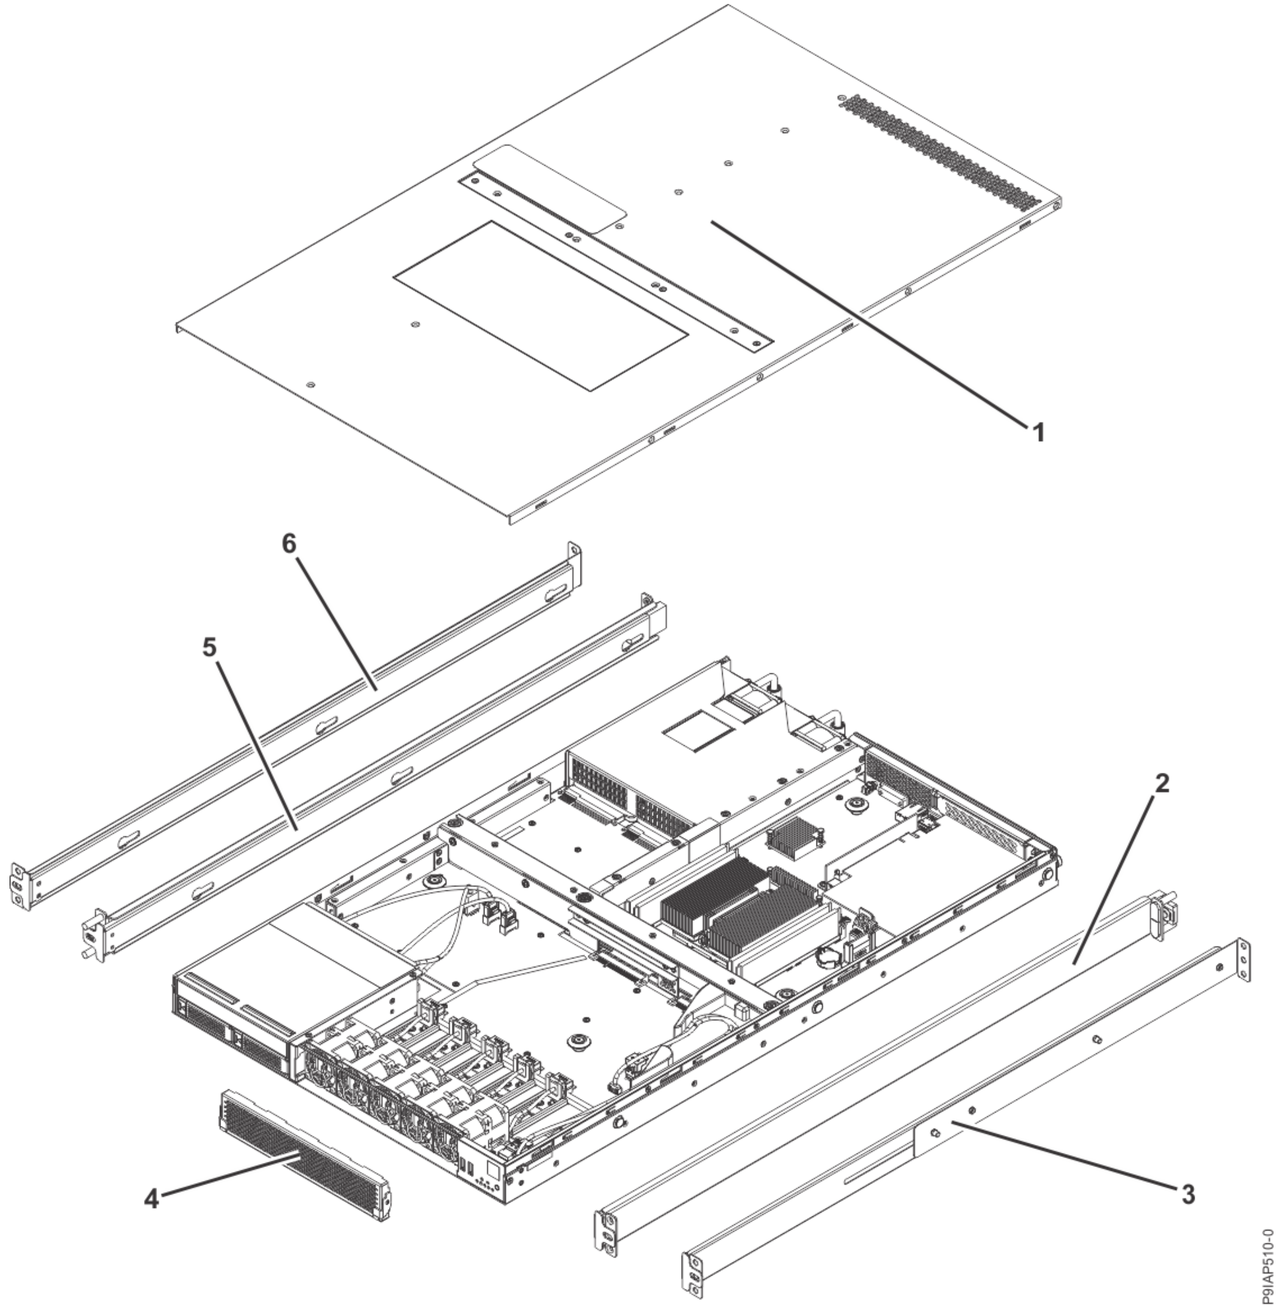

Figure 6. Rack final assembly

Table 10. Rack final assembly part numbers

Index number	Part number	Units per assembly	Description
1		1	Top cover assembly
2	03GM910	1	Fixed rail kit - contains left rail, right rail, attaching screws, and shipping brackets
3	03GM955	1	Fixed rail kit (adjustable in length) - contains left rail, right rail, and attaching screws
4	03GM764	1	Front bezel

Table 10. Rack final assembly part numbers (continued)			
Index number	**Part number**	**Units per assembly**	**Description**
5	03GM910	1	Fixed rail kit - contains left rail, right rail, attaching screws, and shipping brackets
6	03GM955	1	Fixed rail kit (adjustable in length) - contains left rail, right rail, and attaching screws

System parts

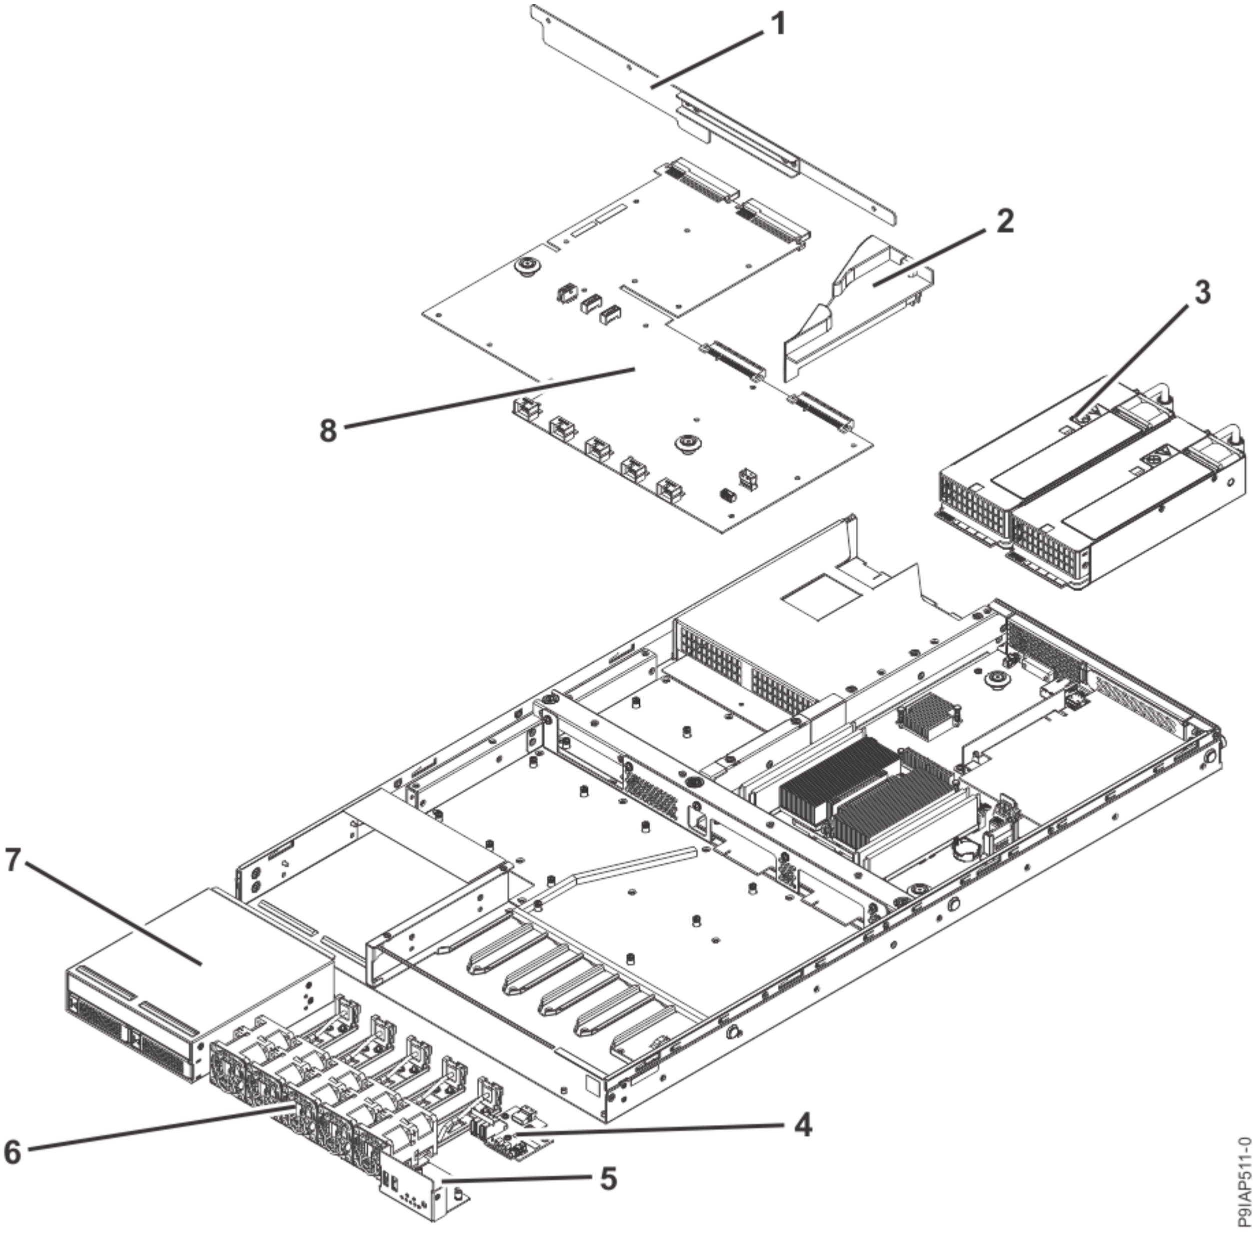

Figure 7. System parts

Table 11. System parts

Index numbe r	Part number	Units per assembly	Description
1	03GM757	1	Air baffle (left)
2	03GM759	1	Air baffle (right)
3	03FP372	2	Power supply
4	02WF445	1	Control panel card
5		1	Control panel cover
6	03GM774	5	Fan
7	03GM776	1	Drive holder
	03GM792	2	1.8 TB 2.5 inch SAS disk drive (includes drive and carrier)
	03GM803	2	Drive carrier
8	02WF441	1	Power distribution board

Additional system parts

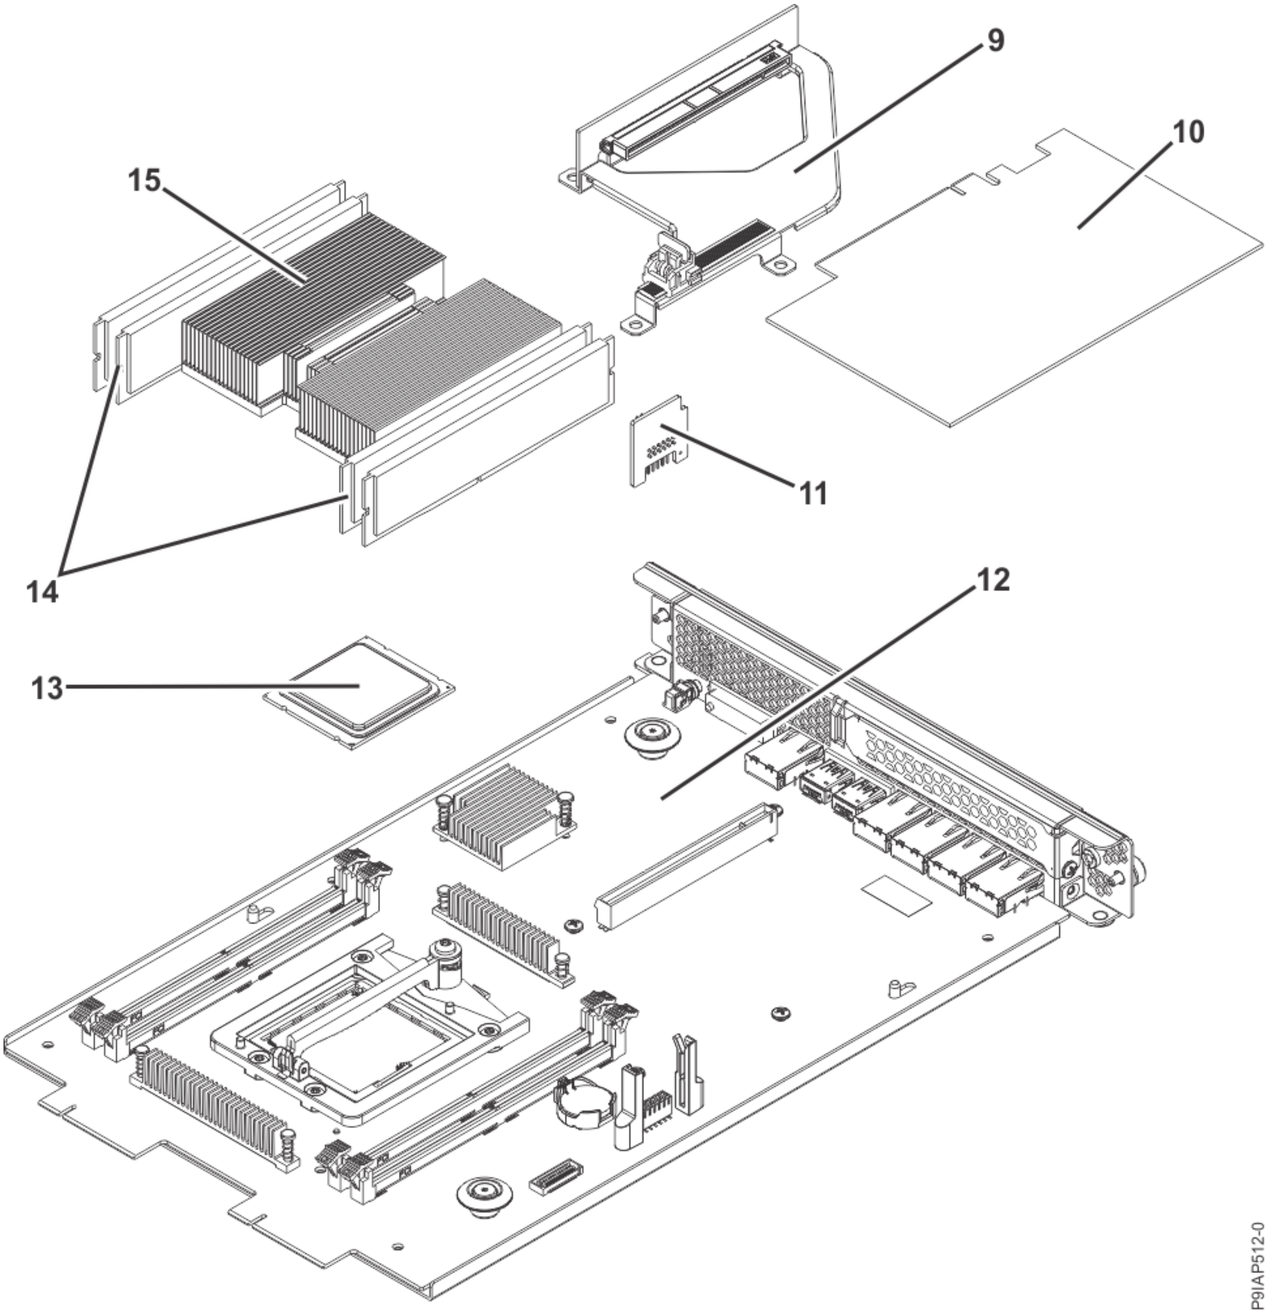

Figure 8. Additional system parts

Table 12. Additional system parts

Index numbe r	Part number	Units per assembly	Description
9	02WF443	1	PCIe riser
10	02JD569	1	PCIe2 2-port 10 GbE BaseT RJ45 adapter
11	00VK865	1	Trusted platform module

Table 12. Additional system parts (continued)			
Index number	**Part number**	**Units per assembly**	**Description**
12	02WF439	1	System backplane kit (includes time-of-day battery, PCIe riser, PCIe tailstock filler, system processor module removal tool, and thermal interface material)
13	03GM808	1	6 core 3.0 GHz system processor module kit (includes system processor module, thermal interface material, tweezers, and system processor module removal tool)
14	78P6722	4	16 GB 2RX4 DDR4 IS RDIMM (Micron Technology, Inc.)
	78P4197	4	16 GB 2RX4 DDR4 IS RDIMM (Samsung Electronics Co., Ltd. or SK hynix, Inc.)
	78P4198	4	32 GB 2RX4 DDR4 IS RDIMM (Micron Technology, Inc., Samsung Electronics Co., Ltd., or SK hynix, Inc.)
15	03GM989	1	Heat sink kit (includes heat sink and thermal interface material)

Table 13. Miscellaneous parts	
Description	**Part number**
USB cable	03GM778
Control panel cable	03GM780
Drive signal cable	03GM782
Drive power cable	03GM784
CR2032 Lithium time-of-day battery	00RY543
Chassis crossbar	03GM755

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details/us/en/ in the section entitled "Cookies, Web Beacons and Other Technologies".

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Electronic emission notices

Class A Notices

The following Class A statements apply to the IBM servers that contain the POWER9 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

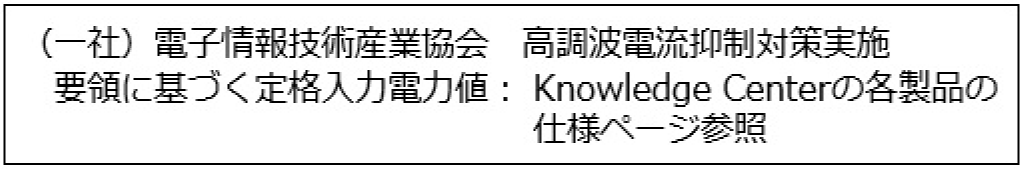

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

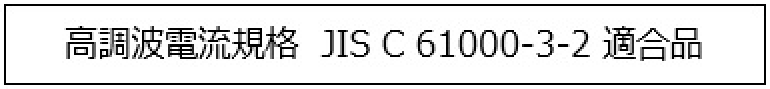

高調波電流規格　JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

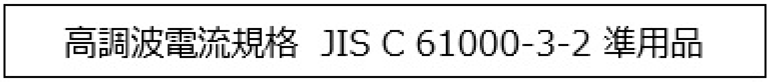

高調波電流規格　JIS C 61000-3-2 準用品

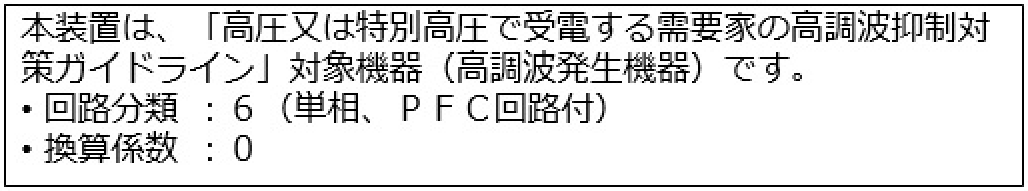

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類　：６（単相、ＰＦＣ回路付）
- 換算係数　：０

This statement applies to products greater than 20 A per phase, three-phase.

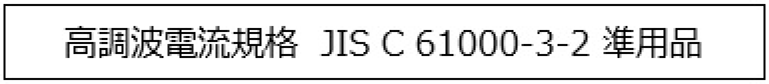

高調波電流規格　JIS C 61000-3-2 準用品

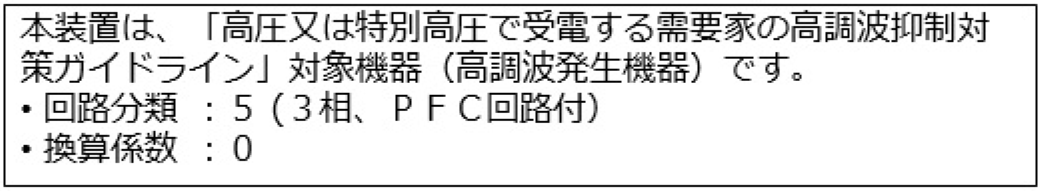

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類　：５（３相、ＰＦＣ回路付）
- 換算係数　：０

Japan Voluntary Control Council for Interference (VCCI) Notice

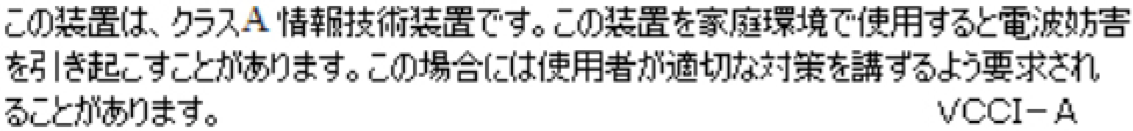

この装置は、クラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI－A

Korea Notice

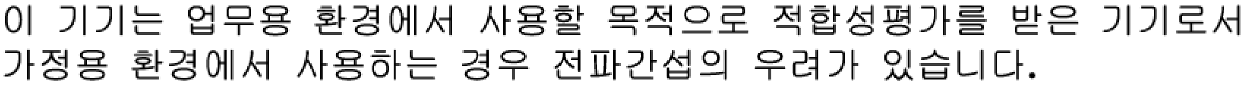

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

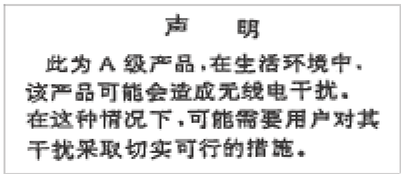

声　明

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰。
在这种情况下，可能需要用户对其
干扰采取切实可行的措施。

Russia Notice

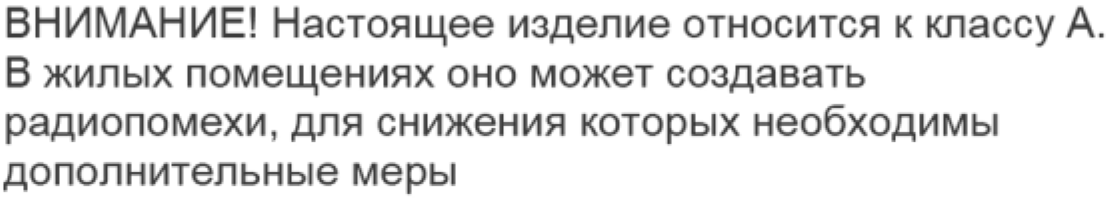

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

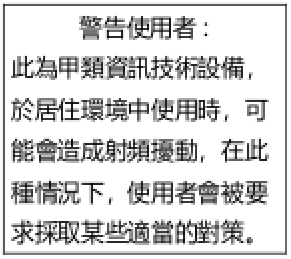

警告使用者：
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

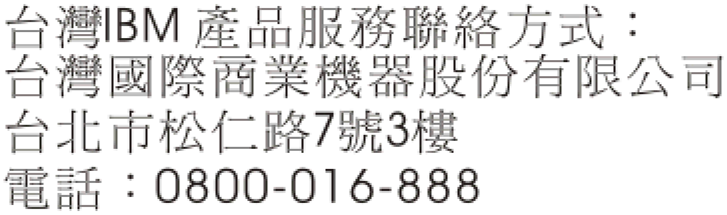

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany

Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

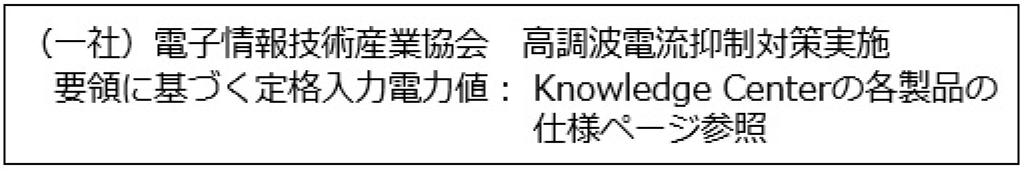

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

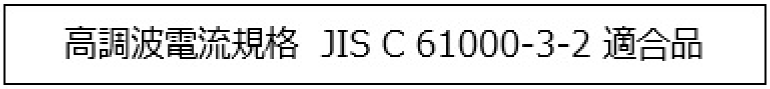

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

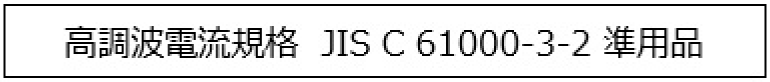

高調波電流規格 JIS C 61000-3-2 準用品

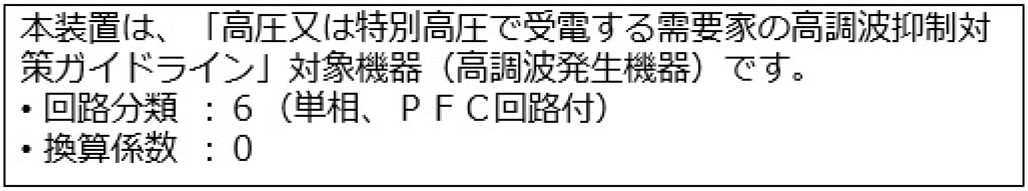

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

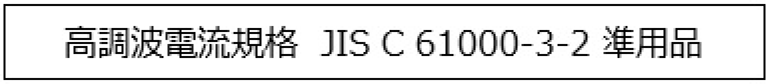

高調波電流規格 JIS C 61000-3-2 準用品

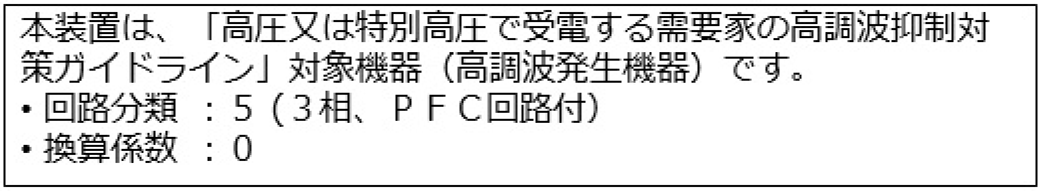

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

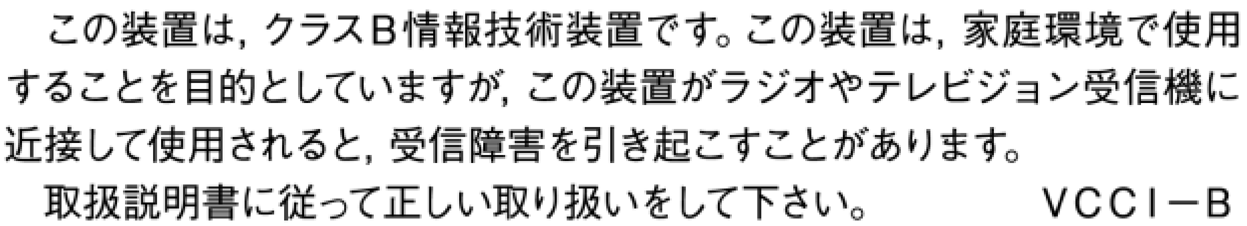

この装置は，クラスＢ情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。　ＶＣＣＩ－Ｂ

Taiwan Notice

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Beginning troubleshooting and problem analysis

Note

Before using this information and the product it supports, read the information in "Safety notices" on page v, "Notices" on page 73, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Power® servers that contain the Power10 processor and to all associated models.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

⚠ **DANGER:** When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.

- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- 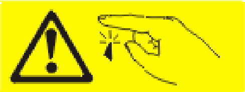 Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).

- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.

  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer's DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

⚠ **CAUTION:**

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position (compliance ID RACK-001 or 22U (compliance ID RR001) and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.

  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U (compliance ID RACK-001 or 22U (compliance ID RR001) level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

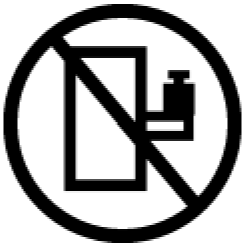

⚠ **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.

- Before extending the rack to the installation position, read the installation instructions.
- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

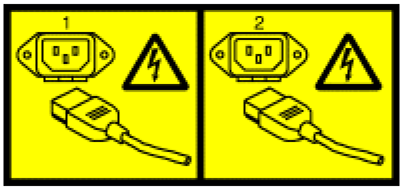

or

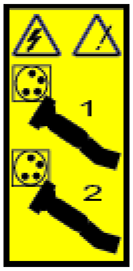

or

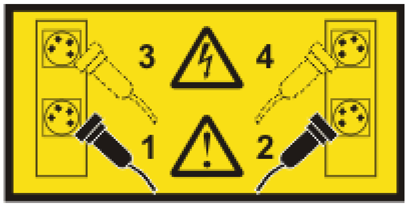

or

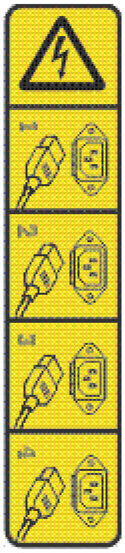

or

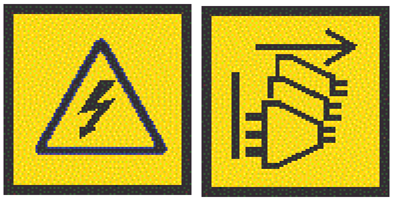

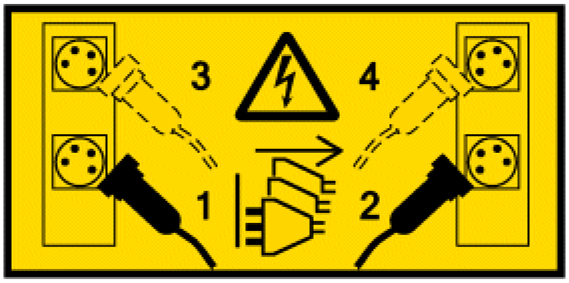

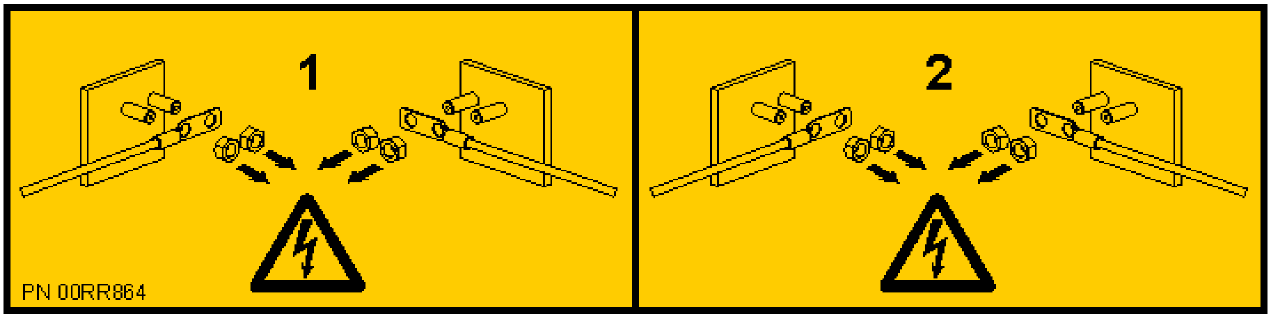

DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

CAUTION: A hot surface nearby. (L007)

(L008)

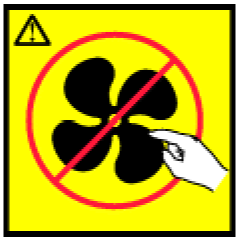

CAUTION: Hazardous moving parts nearby. (L008)

(L018)

or

CAUTION: High levels of acoustical noise are (or could be under certain circumstances) present. Use approved hearing protection and/ or provide mitigation or limit exposure. (L018)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

⚠ **CAUTION:** This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

⚠ **CAUTION:** Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

⚠ **CAUTION:** This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

⚠ **CAUTION:** Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

⚠ **CAUTION:** The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approvedapproved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)(C003a)

⚠ **CAUTION:** Regarding IBM providedprovided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.

- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four (4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.
- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.
- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intra-building ports of this equipment are suitable for connection to intra-building or unexposed wiring or cabling only. The intra-building ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The AC-powered system does not require the use of an external surge protection device (SPD).

The DC-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The DC-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Beginning troubleshooting and problem analysis

This information provides a starting point for analyzing problems.

This information is the starting point for diagnosing and repairing servers. From this point, you are guided to the appropriate information to help you diagnose server problems, determine the appropriate repair action, and then perform the necessary steps to repair the server. A system attention light, an enclosure fault light, or a system information light indicates there is a serviceable event (an SRC in the control panel or in one of the serviceable event views) on the system. This information guides you through finding the serviceable event.

Beginning problem analysis

You can use problem analysis to gather information that helps you determine the nature of a problem encountered on your system. This information is used to determine if you can resolve the problem yourself or to gather sufficient information to communicate with a service provider and quickly determine the service action that needs to be taken.

If you are using this information because of a problem with your Hardware Management Console (HMC), see Managing the HMC.

To begin analyzing the problem, complete the following steps:

1. Do you have a direct indication of a hardware error (such as an automated email that notified you of a hardware error or a fault indicator on a system unit or expansion unit)?

   - **Yes:** Continue with the next step.
   - **No:** Go to "Detecting problems" on page 54.

2. How do you manage the system that is failing? If you do not know how the failing system is managed, ask the system administrator.

System management	Problem analysis
Hardware Management Console (HMC)	Go to the section "Hardware Management Console (HMC) problem analysis" on page 1.
Operating system (AIX®, Linux®, or IBM i)	Go to the problem analysis topic for your operating system. • If you are having a problem with an AIX or Linux system unit, go to "AIX and Linux problem analysis" on page 3. • If you are having a problem with an IBM i system unit, go to " IBM i problem analysis" on page 6.

Hardware Management Console (HMC) problem analysis

To perform beginning problem analysis on a system that is managed by Hardware Management Console (HMC), complete the following steps:

1. Is the management console functional and connected to the hardware?

   - **Yes:** Continue with the next step.
   - **No:** Start the management console and attach it to the system unit. Then return here and continue with the next step.

2. On the management console that is used to manage the system unit, complete the following steps:

   Note: If you are unable to locate the reported problem, and there is more than one open problem near the time of the reported failure, use the earliest problem in the log.

   a. In the navigation area, click **Serviceability**, and then click **Serviceable Events Manager**. The Manage Serviceable Events window is displayed.
   b. In the Event Criteria area, for **Serviceable Event Status**, select **Open**. For all other criteria, select **ALL**, then click **OK**.

   Scroll through the log and verify that there is a problem with the status of Open to correspond with the failure.

   Do you find a serviceable event, or an open problem near the time of the failure?

   - **Yes:** Continue with the next step.
   - **No:** Contact your hardware service provider. **This ends the procedure.**

3. The reference code description might provide information or an action that you can take to correct the failure.

   Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action at this time.

   For more information about reference codes, see Reference codes.

   Was there a reference code description that enabled you to resolve the problem?

   - **Yes: This ends the procedure.**
   - **No:** Continue with the next step.

4. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

   - If a FRU location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
   - If an isolation procedure is listed for the reference code in the reference code lookup information, include the isolation procedure as a corrective action even if it is not listed in the serviceable event view or control panel.
   - If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

   From the Repair Serviceable Event window, complete the following steps:

   a. Record the problem management record (PMR) number for the problem if one is listed.
   b. Select the serviceable event from the list.
   c. Click **Selected and View Details**.
   d. In the Serviceable Event Details page, locate details such as the reference code and FRU list and record this information.
   e. If a Problem Management Hardware (PMH) number was found for the problem on the Serviceable Event Overview panel, the problem has already been reported. If there was no PMH number for the problem, contact your service provider.

   This ends the procedure.

AIX and Linux problem analysis

You can use this procedure to find information about a problem with your server hardware when service is managed by the AIX or Linux operating system.

Remember the following points while troubleshooting problems:

- Has an external power outage or momentary power loss occurred?
- Has the hardware configuration changed?
- Has system software been added?
- Have any new programs or program updates (including PTFs) been installed recently?

Before you use this procedure, ensure that you completed the steps in "Beginning problem analysis" on page 1.

After you review these considerations, complete the following steps:

1. Is the operating system operational?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "11" on page 5.

2. Are any messages (for example, a device is not available or reporting errors) related to this problem displayed on the system console or sent to you in email that provides a reference code?

   Note: A reference code can be an 8 character system reference code (SRC) or a service request number (SRN) of 5, 6, or 7 characters, with or without a hyphen.

   - **Yes:** Continue with the next step.
   - **No:** Go to step "4" on page 3.

3. The reference code description might provide information or an action that you can take to correct the failure.

   Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

   For more information about reference codes, see Reference codes.

   If the reference code description provides information to resolve the problem without replacing FRUs in the failing item list, complete the steps.

   Were you able to resolve the problem?

   - **Yes: This ends the procedure.**
   - **No:** Continue with the next step.

4. Are you running the Linux operating system?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "6" on page 4.

5. To locate the error information in a system or logical partition that is running the Linux operating system, complete the following steps:

   Note: Before you proceed with this step, ensure that the diagnostics package is installed on the system.

a. Log in as root user.
b. At the command line, type `grep RTAS /var/log/platform` and press **Enter**.
c. Look for the most recent entry that contains a reference code.

Continue with step “8” on page 4.

6. To locate the error information in a system or logical partition that is running the AIX operating system, complete the following steps:

   a. Log in to the AIX operating system as root user, or use CE login. If you need help, contact the system administrator.
   b. Type `diag` to load the diagnostic controller, and display the online diagnostic menus.
   c. From the Function selection menu, select **Task selection**.
   d. From the Task selection list menu, select **Display previous diagnostic results**.
   e. From the Previous diagnostic results menu, select **Display diagnostic log summary**.

   Continue with the next step.

7. A display diagnostic log is shown with a time ordered table of events from the error log.

   Look in the T column for the most recent entry that has an S entry. Press **Enter** to select the row in the table and then select **Commit**.

   The details of this entry from the table are shown. Look for the SRN entry near the end of the entry and record the information that is shown.

   Continue with the next step.

8. Do you find a serviceable event or an open problem near the time of the failure?

   - **Yes:** Continue with the next step.
   - **No:** Contact your hardware service provider. **This ends the procedure.**

9. The reference code description might provide information or an action that you can take to correct the failure.

   Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

   For more information about reference codes, see Reference codes.

   Was there a reference code description that helped you to resolve the problem?

   - **Yes: This ends the procedure.**
   - **No:** Continue with the next step.

10. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

- If a field-replaceable unit (FRU) location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
- If an isolation procedure is listed for the reference code in the reference code lookup information, include it as a corrective action even if it is not listed in the serviceable event view or control panel.
- If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

From the Error Event Log view, complete the following steps:

a. Record the reference code.
b. Record the error details.
c. Contact your service provider.

This ends the procedure.

---

11. Details about errors that occur when the operating system is not running or when the operating system is now not accessible can be found in the control panel or in the Advanced System Management Interface (ASMI).

    Do you choose to look for error details by using ASMI?

    - **Yes:** Go to step "13" on page 5.
    - **No:** Continue with the next step.

---

12. At the control panel, complete the following steps.

    a. Press the increment or decrement button until the number 11 is displayed in the upper-left corner of the display.
    b. Press **Enter** to display the contents of function 11.
    c. Look for a reference code in the upper-right corner.

    Is a reference code displayed on the control panel in function 11?

    - **Yes:** Go to step "14" on page 5.
    - **No:** Contact your hardware service provider. **This ends the procedure.**

---

13. On the console that is connected to the ASMI, complete the following steps.

    Note: If you are unable to locate the reported problem, and there is more than one open problem near the time of the reported failure, use the earliest problem in the log.

    a. Log in with a user ID that has an authority level as general, administrator, or authorized service provider.
    b. In the navigation area, expand **System Service Aids** and click **Error/Event Logs**. If log entries exist, a list of error and event log entries is displayed in a summary view.
    c. Scroll through the log under **Serviceable Customer Attention Events** and verify that there is a problem to correspond with the failure.

    For information about the ASMI, see Managing the Advanced System Management Interface.

    Do you find a serviceable event, or an open problem near the time of the failure?

    - **Yes:** Continue with the next step.
    - **No:** Contact your hardware service provider. **This ends the procedure.**

---

14. The reference code description might provide information or an action that you can take to correct the failure.

Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

For more information about reference codes, see Reference codes.

Was there a reference code description that helped you to resolve the problem?

- **Yes: This ends the procedure.**
- **No:** Continue with the next step.

15. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

    - If a field-replaceable unit (FRU) location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
    - If an isolation procedure is listed for the reference code in the reference code lookup information, include the isolation procedure as a corrective action even if it is not listed in the serviceable event view or control panel.
    - If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

    To find error details on the control panel, complete the following steps:

    a. Press **Enter** to display the contents of function 14. If data is available in function 14, the reference code has a FRU list.
    b. Record the information in functions 11 through 20 on the control panel.
    c. Contact your service provider and report the reference code and other information.

    To find error details on the ASMI, complete the following steps from the Error Event Log view:

    a. Record the reference code.
    b. Select the corresponding check box on the log and click Show details.
    c. Record the error details.
    d. Contact your service provider.

    This ends the procedure.

IBM i problem analysis

You can use this procedure to find information about a problem with your server hardware when service is managed by the IBM i operating system.

If you experience a problem with your system or logical partition, try to gather more information about the problem to either solve it, or to help your next level of support or your hardware service provider to solve it more quickly and accurately.

This procedure refers to the IBM i control language (CL) commands that provide a flexible means of entering commands on the IBM i logical partition or system. You can use CL commands to control most of the IBM i functions by entering them from either the character-based interface or the IBM Navigator for i Web console. While the CL commands might be unfamiliar at first, the commands follow a consistent syntax, and IBM i includes many features to help you use them easily. The Programming navigation category in IBM i Documentation includes a complete CL reference and a CL Finder to look up specific CL commands.

Remember the following points while troubleshooting problems:

- Has an external power outage or momentary power loss occurred?

- Has the hardware configuration changed?
- Has system software been added?
- Have any new programs or program updates (including PTFs) been installed recently?

To make sure that your IBM software was correctly installed, use the Check Product Option (CHKPRDOPT) command.

- Have any system values changed?
- Has any system tuning been done?

Before you use this procedure, ensure that you completed the steps in "Beginning problem analysis" on page 1.

After you review these considerations, follow these steps:

1. Is the IBM i operating system up and running?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "19" on page 9.

2. Are you experiencing problems with the Operations Console?

   - **Yes:** See Troubleshooting Operations Console.
   - **No:** Continue with the next step.

3. Does the console show a Main Storage Dump Manager display?

   - **Yes:** Go to Copying a dump.
   - **No:** Continue with the next step.

4. Is the console that was in use when the problem occurred (or any console) operational?

   Note: The console is operational if a sign-on display or a command line is present. If another console is operational, use it to resolve the problem.

   - **Yes:** Continue with the next step.
   - **No:** See the Troubleshooting navigation category in IBM i Documentation.

5. Is a message related to this problem shown on the console?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "10" on page 8.

6. Is this a system operator message?

   Note: It is a system operator message if the display indicates that the message is in the QSYSOPR message queue. Critical messages can be found in the QSYSMSG message queue. For more information, see the *Create message queue QSYSMSG for severe messages* topic in the Troubleshooting navigation category of IBM i Documentation.

   - **Yes:** Continue with the next step.
   - **No:** Go to step "8" on page 8.

7. Is the system operator message highlighted, or does it have an asterisk (*) next to it?

   - **Yes:** Go to step "17" on page 9.
   - **No:** Go to step "12" on page 8.

8. Move the cursor to the message line and press F1 (Help). Does the Additional Message Information display appear?

   - **Yes:** Continue with the next step.
   - **No:** Go to step "10" on page 8.

---

9. Record the additional message information on the appropriate problem reporting form. For details, see "Problem reporting form" on page 14.

   Follow the recovery instructions on the Additional Message Information display.

   Did this solve the problem?

   - **Yes: This ends the procedure**.
   - **No:** Continue with the next step.

---

10. To display system operator messages, type `dspmsg qsysopr` on any command line and then press **Enter**.

    Did you find a message that is highlighted or has an asterisk (*) next to it?

    - **Yes:** Go to step "17" on page 9.
    - **No:** Continue with the next step.

    Note: The message monitor in the IBM Navigator for i Web console can also inform you when a problem has developed. For details, see the *Scenario: Message monitor topic in the Systems Management navigation category* of IBM i Documentation.

---

11. Did you find a message with a date or time that is at or near the time the problem occurred?

    Note: Move the cursor to the message line and press F1 (Help) to determine the time that a message occurred. If the problem is shown to affect only one console, you might be able to use information from the JOB menu to diagnose and solve the problem. To find this menu, type **GO JOB** and press **Enter** on any command line.

    - **Yes:** Continue with the next step.
    - **No:** Go to step "14" on page 9.

---

12. Complete the following steps:

    a. Move the cursor to the message line and press F1 (Help) to display additional information about the message.
    b. Record the additional message information on the appropriate problem reporting form. For details, see "Problem reporting form" on page 14.
    c. Follow any recovery instructions that are shown.

    Did this solve the problem?

    - **Yes: This ends the procedure**.
    - **No:** Continue with the next step.

---

13. Did the message information indicate to look for additional messages in the system operator message queue (QSYSOPR)?

    - **Yes:** Press F12 (Cancel) to return to the list of messages and look for other related messages. Then, return to step "10" on page 8.
    - **No:** Continue with the next step.

---

14. Do you know which input/output device is causing the problem?

    - **Yes:** Continue with step "16" on page 9.
    - **No:** Continue with the next step.

15. If you do not know which input/output device is causing the problem, describe the problems that you observed by completing the following steps:

    a. Type `GO USERHELP` on any command line and then press **Enter**.
    b. Select option 10 (Save information to help resolve a problem).
    c. Type a brief description of the problem and then press **Enter**. If you specify the default **Y** in the **Enter notes about problem** field, you can enter more text to describe your problem.
    d. Report the problem to your hardware service provider.

16. Complete the following steps:

    a. Type ANZPRB on the command line and then press **Enter**. For details, see *Using the Analyze Problem (ANZPRB) command* in the Troubleshooting navigation category in IBM i Documentation.
    b. Contact your next level of support. **This ends the procedure**.

    Note: To describe your problem in greater detail, see *Using the Analyze Problem (ANZPRB) command* in the Troubleshooting navigation category in IBM i Documentation. This command can also run a test to further isolate the problem.

17. Complete the following steps:

    a. Move the cursor to the message line and press F1 (Help) to display additional information about the message.
    b. Press F14, or use the Work with Problem (WRKPRB) command. For details, see *Work with Problem (WRKPRB)* in the Troubleshooting navigation category in IBM i Documentation.
    c. If this does not solve the problem, see the IBM i server or IBM i partition symptoms.

18. Choose from the following options:

    - If reference codes appear on the control panel or the management console, record them. Then, go to Reference codes to see if additional details are available for the code you received.
    - If no reference codes appear on the control panel or the management console, a serviceable event is indicated by a message in the problem log. Use the WRKPRB command. For details, see *Work with Problem (WRKPRB)* in the Troubleshooting navigation category in IBM i Documentation.

19. Details about errors that occur when IBM i is not running or when IBM i is now not accessible can be found in the control panel or in the Advanced System Management Interface (ASMI).

    Do you choose to look for error details by using ASMI?

    - **Yes:** Go to step "21" on page 10.
    - **No:** Continue with the next step.

20. At the control panel, complete the following steps.

a. Press the increment or decrement button until the number 11 is displayed in the upper-left corner of the display.
b. Press **Enter** to display the contents of function 11.
c. Look for a reference code in the upper-right corner.

Is there a reference code displayed on the control panel in function 11?

- **Yes:** Go to step "22" on page 10.
- **No:** Contact your hardware service provider. **This ends the procedure.**

---

21. On the console that is connected to the ASMI, complete the following steps.

Note: If you are unable to locate the reported problem, and there is more than one open problem near the time of the reported failure, use the earliest problem in the log.

a. Log in with a user ID that has an authority level as general, administrator, or authorized service provider.
b. In the navigation area, expand **System Service Aids** and click **Error/Event Logs**. If log entries exist, a list of error and event log entries is displayed in a summary view.
c. Scroll through the log under **Serviceable Customer Attention Events** and verify that there is a problem to correspond with the failure.

For more detailed information on the ASMI, see Managing the Advanced System Management Interface.

Do you find a serviceable event, or an open problem near the time of the failure?

- **Yes:** Continue with the next step.
- **No:** Contact your hardware service provider. **This ends the procedure.**

---

22. The reference code description might provide information or an action that you can take to correct the failure.

Use the search function of IBM Knowledge Center to find the reference code details. The search function is located in the upper-left corner of IBM Knowledge Center. Read the reference code description and return here. Do not take any other action now.

For more information about reference codes, see Reference codes.

Was there a reference code description that helped you to resolve the problem?

- **Yes: This ends the procedure.**
- **No:** Continue with the next step.

---

23. Service is required to resolve the error. Collect as much error data as possible and record it. You and your service provider will develop a corrective action to resolve the problem based on the following guidelines:

- If a field-replaceable unit (FRU) location code is provided in the serviceable event view or control panel, use that location to determine which FRU to replace.
- If an isolation procedure is listed for the reference code in the reference code lookup information, include the isolation procedure as a corrective action even if it is not listed in the serviceable event view or control panel.
- If any FRUs are marked for block replacement, replace all FRUs in the block replacement group at the same time.

To find error details on the control panel, complete the following steps:

a. Press **Enter** to display the contents of function 14. If data is available in function 14, the reference code has a FRU list.
b. Record the information in functions 11 through 20 on the control panel.
c. Contact your service provider and report the reference code and other information.

To find error details on the ASMI, complete the following steps from the Error Event Log view:

a. Record the reference code.
b. Select the corresponding check box on the log and click Show details.
c. Record the error details.
d. Contact your service provider.

This ends the procedure.

Light path diagnostics

Light path diagnostics is a simplified approach for repair actions that provides fault indicators to identify parts that need to be replaced.

Light path diagnostics is a system of light-emitting diodes (LEDs) on the control panel and on various internal components. When an error occurs, LEDs are lit throughout the system to help identify the source of the error.

With light path diagnostics, the fault LED for the FRUs to be replaced will be active when the unit is powered on. The failing FRUs can be attached to another FRU, which must be first removed to access the failing FRUs. For those cases, light path diagnostics provides a blue switch on the FRU that has to be removed first. When the first FRU is removed, you can press and hold the light path diagnostics switch to light the LEDs and locate the failing part. In most of the situations, the switch will have enough power stored to activate the LEDs for two hours after the unit has been powered off. However, this can vary significantly and therefore the switch should be used as soon as possible. The amber LEDs can normally be kept active for 30 seconds, however, this can also vary. Associated with the light path diagnostics switch is a green LED that will be activated when the switch is used and there is enough power stored to activate the amber LEDs. If the green LED does not activate when the switch is pressed then there is not enough power remaining to activate any amber LEDs on that FRU. If that happens then light path diagnostics and FRU identify function cannot be used for replacing the failing FRUs. Perform the repair action using the location codes in the error log or if determined by problem analysis as if the unit did not have light path diagnostics and did not have functioning identify LEDs.

At the core of light path diagnostics is a set of fault indicators that are implemented as amber LEDs. These LEDs provide a way for the diagnostics to identify which field-replaceable unit (FRU) needs to be replaced. Service labels, color-coded touch points for the FRUs, and a no tools required design for FRU removal and installation are all elements of light path diagnostics.

With light path diagnostics, at the same time that the diagnostics create an error log for the problem, it also activates the fault indicator when a FRU has a failed or failing component. This includes predictive failure analysis (PFA). The FRU fault LED is turned on solid (not flashing). Whenever a fault indicator is activated the enclosure's external Fault indicator on the operator panel is also turned on solid. The enclosure fault indicator on the panel means that inside the unit one or more FRU fault indicators is on.

The error log shows the part number and location code of the FRU that must be replaced along with other FRUs or procedures to follow if replacing the first FRU does not resolve the problem.

If the diagnostics determine that the problem is firmware related, configuration related, or not isolated to a specific FRU then no fault indicator is activated. For these kinds of problems the amber System Info indicator on the operator panel is activated. The error log shows the procedures to follow and the possible FRUs that could be causing the problem.

During the repair action, the service label on the service access cover shows the FRUs and the steps required to remove or install the FRUs. Therefore, the basic flow of the repair is that the LEDs show which part to replace, the color-coded touch points indicate if the unit must be powered off to remove or install the part, and the service label shows the steps needed on the touch points. The FRU fault LED is not an indication that the FRU is ready to be replaced. To replace the FRU, some preparation steps might be necessary, such as removing the resource from use or powering off the unit. The service label and the touch point colors give the initial guidance for FRU removal.

When a FRU has been replaced, the fault indicator automatically turns off either when the new FRU is installed, or when the power is restored to the new FRU. This automatic shut off might take several seconds to a minute as the new FRU is powered on, brought online and tested by the system. When there are no more FRU fault indicators on in an enclosure then the enclosure's Fault indicator on the operator panel turns off automatically.

In addition to the fault indicators, there are also amber identify indicators for each FRU. The identify indicators flash when activated. Identify indicators are used to help a servicer identify where a location is. The location may be occupied or empty. The servicer can turn them on and off from a user interface either during a repair action or during the installation of new parts or when removing parts. The identify indicator visually confirms where a location code is. Whenever an identify indicator is activated, the enclosure's blue *locate* or *beacon* LED on the operator panel is also activated (flashing).

The same amber LED on a FRU can be used for both fault and identify indications. Whenever the LED is on solid for a fault, the LED switches to flashing when the FRU identify function is turned on. When identify function is turned off, the LED returns to fault (solid on) if that was the previous state of the LED.

Replacing FRUs by using enclosure fault indicators

After you obtain a replacement part, use this procedure to identify the location of the part that needs to be replaced.

About this task

To identify and locate the part that requires replacement, complete the following steps.

Procedure

1. Before you move the unit into the service position, refer to the service label. It might be necessary to identify and remove cables that are attached to the FRU you are about to exchange. Use the FRU location code and the service label to determine whether there are any actions before you move the unit into the service position. Complete those actions and return to the next step in this procedure.
2. Identify the unit with the active enclosure fault indicator. Use the service label that is affixed to the service access cover and the amber light-emitting diode (LED) on the FRU to find the failing FRU. Move the unit into the service position, but do not remove the service access cover.
   - If the unit is rack-mounted, the service label is visible on the service access cover. Continue with the next step.
   - If the unit is a stand-alone system, the exterior cover must be removed to view the service label.
3. Using the service label, determine whether the FRU you are replacing can be exchanged without removing the service access cover. Is the FRU fault LED visible externally and active (on solid, not flashing) and does the service label show that the service access cover does not need to be removed to exchange the FRU? (Choose No if you are unsure.)

Note: If you used the identify function in a user interface to help locate the FRU, then the amber LED is flashing. Otherwise, the amber LED is solid (not flashing).

- **Yes:** Go to step 6.
- **No:** To identify which FRU to exchange, you must remove the service access cover and locate the FRU that has an active FRU fault indicator (amber LED on). Continue with the next step.

4. Remove the service access cover and locate the FRU that has an active fault indicator (amber LED on, not flashing). Use the following table to determine whether you must power off the unit before you remove the cover.

   Note: You can remove the service access cover while the unit is powered on.
5. Search for the FRU to be replaced by locating the active amber LED.

   Notes:

   - If you used the identify function in a user interface to help locate the FRU, then the amber LED is flashing. Otherwise, the amber LED is solid (not flashing).
   - Some FRUs might be an integral part of another FRU. This might make it difficult to see the FRU that you need to exchange or the amber LED that designates the FRU that needs to be exchanged. If so, remove all FRUs that are associated with the failing FRU.

   Do you need to remove another FRU to replace the FRU designated by the amber LED?

   - **Yes:** Go to step 9.
   - **No:** Continue with the next step.
6. For the FRU you located with its fault LED active, was it replaced for this problem or service action?

   - **Yes:** The FRU replaced for the original problem did not resolve the problem. Go back to the serviceable event for the original problem and address the remaining FRUs that are listed.

     Note: If the fault indicator for the replaced FRU is turned on, use the Advanced System Management Interface (ASMI) to turn off the fault indicator.

     This ends the procedure.
   - **No:** Continue with the next step.
7. For the FRU you located with the active fault LED, compare the location code that you recorded for the replacement FRU of the problem you are working on with the location code of the active fault indicator. If they do not match, you are working a problem from the log that is different from the problem that activated the fault indicators. Do the location codes match?

   - **Yes:** You are working the same problem that activated the fault indicators. Continue with the next step.
   - **No:** If you have the correct replacement FRU for where the fault indicator is active, you can continue with this repair action. When you replace the FRU, record the location codes of the active fault indicators for use later to identify which problem to close when the repair is complete, then continue with the next step. Otherwise, contact your service provider to obtain the replacement part for the FRU with an active fault indicator and begin problem analysis again. **This ends the procedure.**
8. If you have not already done so, record the location of the FRU you are about to exchange either by where the service label shows the FRU or where it is in the unit by the location labeling. For information on part locations and location codes, see Part locations and location codes. In the locations and addresses information for your system, locate the FRU and the corresponding FRU exchange procedure. The exchange procedure provides the steps necessary to exchange the FRU. If the FRU can be replaced while the unit is powered on, the exchange procedure provides that option and the necessary instructions. If the enclosure fault indicator turns on again within a few minutes of completing the replacement and returning to normal use of the unit, then begin problem analysis again. Otherwise, close the problem. **This ends the procedure.**

9. If you have not already done so, record the location of the FRU that you plan to remove either by where the service label shows the FRU or by the location labeling where it is in the unit. For information on part locations and location codes, see Part locations and location codes.

   In the locations and addresses information, locate the FRU and the corresponding FRU exchange procedure. The exchange procedure provides the steps necessary to remove this FRU. If the FRU can be removed while the unit is powered on, the exchange procedure provides that option and the necessary instructions. When you remove this FRU, the fault indicator for the associated FRU you are exchanging turns off. This FRU has an LED activation button that you can press that powers the amber indicator of the FRU you are exchanging. Use the button to locate the FRU you are exchanging. Go to step 8.

   Note: If the button's green LED does not activate, there is not enough charge in the switch to activate the amber fault LED. To identify the failing FRU, use the FRU location code either from the error log or as determined by problem analysis.

10. For the FRU you located that has its fault LED active, were any of them replaced for this problem or service action?

    - **Yes:** The FRU replaced for the original problem did not resolve the problem. Go back to the serviceable event for the original problem and work the remaining FRUs listed. Use ASMI to turn off the fault indicator for the FRU. **This ends the procedure.**
    - **No:** Use the information on the service label to exchange the FRU. When the FRU is exchanged, use the service label to guide you in reassembling the unit. Power on the unit. The FRU fault indicator is turned off during the power-on process if it was not already turned off. If the enclosure fault indicator turns on again within a few minutes of powering on the unit, then begin problem analysis again. Otherwise, close this problem. **This ends the procedure.**

Problem reporting form

Use the problem reporting form to record information about your server that will assist you in problem analysis.

Collect as much information as possible in the tables below, using either the control panel or the management console to gather the information.

Table 1. Customer, system, and problem information

Customer information and problem description	
Your name	
Telephone number	
IBM customer number, if available	
Date and time that the problem occurred	
Describe the problem	
System Information	
Machine type	
Model	
Serial number	
IPL type	
IPL mode	
Message information	

Table 1. Customer, system, and problem information (continued)

Customer information and problem description	
Message ID	
Message text	
Service request number (SRN)	
In which mode were IBM hardware diagnostics run?	___ Online or ___ stand-alone ? ___ Service mode or ___ concurrent mode?

Go to the management console or the control panel and indicate whether the following lights are on.

Table 2. Control panel lights

Control panel light	Place a check if light is on
Power On	
System Attention	

Go to the management console or control panel to find and record the values for functions 11 through 20. Use the following grid to record the characters shown on the management console or Function/Data display.

Table 3. Function values

Function	Value
11	__ __
12	__ __
13	__ __
14	__ __
15	__ __
16	__ __
17	__ __
18	__ __
19	__ __
20 (for control panel users)	__ __

Table 3. Function values (continued)

Function	Value
20 (for management console users)	Machine type: Model: Processor feature code: IPL type:

Reference information for problem determination

The problem determination reference information is provided as an additional resource for problem detection and analysis when you are directed here by your service representative.

All repair actions should start with "Beginning problem analysis" on page 1 before you use these tools and techniques for problem determination.

Symptom index

Use this symptom index only when you are guided here by your service representative.

Note: If you were not guided here by your service representative, go to "Beginning problem analysis" on page 1.

Review the symptoms in the left column. Look for the symptom that most closely matches the symptoms on the server that you are troubleshooting. When you find the matching symptom, perform the appropriate action as described in the right column.

Table 4. Determining symptom types

Symptom	What you should do:
You do not have a symptom.	Go to the "Beginning problem analysis" on page 1.
The symptom or problem is on a server or a partition running IBM ian operating system other than AIX or Linux.	Go to "IBM i server or IBM i partition symptoms" on page 16.
The symptom or problem is on a server or a partition running AIX.	Go to "AIX server or AIX partition symptoms" on page 21.
The symptom or problem is on a server or a partition running Linux.	Go to "Linux server or Linux partition symptoms" on page 35.

IBM i server or IBM i partition symptoms

Use the following tables to find the symptom you are experiencing. If you cannot find your symptom, contact your next level of support.

- General symptoms
- Symptoms occurring when the system is not operational
- Symptoms related to a logical partition on a server that has multiple logical partitions
- Obvious physical symptoms
- Time-of-day symptoms

Table 5. General IBM i server or IBM i partition symptoms

Symptom	**Service action**
You have an intermittent problem or you suspect that the problem is intermittent.	Go to "Intermittent problems" on page 61.
DST/SST functions are available on the logical partition console and: • The customer reports reduced system function. • There is a server performance problem. • There are failing, missing, or inoperable server resources.	On most servers with logical partitions, it is common to have one or more missing or non-reporting system bus resource's under Hardware Service Manager (see Hardware Service Manager for more information).
Operations Console, or the remote control panel is not working properly.	Contact Software Support.
The system has a processor or memory problem.	Use the Service action log to check for a reference code or any failing items. See Service action log for instructions, replacing any hardware FRUs if necessary.
The system has detected a bus problem. An SRC of the form B600 69*xx* or B700 69*xx* will be displayed on the control panel or management console.	Go to Service action log.

Table 6. Symptoms occurring when the system is not operational

Symptom	**Service action**
A bouncing or scrolling ball (moving row of dots) remains on the operator panel display, or the operator panel display is filled with dashes or blocks.	Verify that the operator panel connections to the system backplane are connected and properly seated. Also, reseat the Service Processor card. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Operator panel assembly. 2. Service processor.

Table 6. Symptoms occurring when the system is not operational (continued)

Symptom	Service action
You have a blank display on the operator panel. Other LEDs on the operator panel appear to behave normally.	Verify that the operator panel connections to the system backplane are connected and properly seated. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Control (operator) panel assembly. 2. Service processor.
There is an IPL problem, the system attention light is on, and blocks of data appear for 5 seconds at a time before moving to the next block of data for 5 seconds, and so on until 5 seconds of a blank control panel is displayed at which time the cycle repeats.	These blocks of data are functions 11 through 20. The first data block after the blank screen is function 11, the second block is function 12, and so on. Use this information to fill out the Problem reporting forms. Then go to Reference codes.
You have a power problem, the system or an attached unit will not power on or will not power off, or there is a 1*xxx*-*xxxx* reference code.	Go to Power problems.
There is an SRC in function 11.	Look up the reference code (see Reference codes).
There is an IPL problem.	Go to "IPL problems" on page 67.
There is a `Device Not Found` message during an installation from an alternate installation device.	Go to TUPIP06.

Table 7. Symptoms related to a logical partition on a server that has multiple logical partitions

Symptom	Service action
The console is not working for a logical partition.	See the Troubleshooting navigation category in IBM i Documentation.

Table 7. Symptoms related to a logical partition on a server that has multiple logical partitions (continued)

Symptom	Service action
• There is an SRC on the panel of an I/O expansion unit owned by a logical partition. • You suspect a power problem with resources owned by a logical partition. • There is an IPL problem with a logical partition and there is an SRC on the management console. • The logical partition's operations have stopped or the partition is in a loop and there is an SRC on the management console.	• Search Service Focal Point for a serviceable event. • If you do not find a serviceable event in Service Focal Point, then record the partition's SRC from the Operator Panel Values field in the management console. 1. In the navigation area, click **Resources**, and then select **All Systems**. 2. In the content pane, select the required system or click on the server name and then select the required partition. 3. Use that SRC and look up the reference code. For more information, see Reference codes.
The logical partition's console is functioning, but the state of the partition in the management console is "Failed" or "Unit Attn" and there is an SRC.	Use the logical partition's SRC. From the partition's console search for that SRC in the partition's Service Action Log. See Service action log.
There is an IPL problem with a logical partition and there is no SRC displayed in the management console.	Perform the following to look for the panel value for the partition in the management console. 1. In the navigation area, click **Resources**, and then select **All Systems**. 2. In the content pane, select the required system or click on the server name and then select the required partition. 3. In the navigation area, click **Serviceability** > **Reference Code Log** and view the codes. 4. When finished, click **Cancel**. Go to Reference codes. If no reference code can be found, contact your next level of support.
The partition's operations have stopped or the partition is in a loop and there is no SRC displayed on the management console.	Perform function 21 from the management console. If this fails to resolve the problem, contact your next level of support.
One or more of the following was reported: • There is a system reference code or message on the logical partition's console. • The customer reports reduced function in the partition. • There is a logical partition performance problem. • There are failing, missing, or inoperable resources.	From the partition's console search the partition's Service Action Log. Go to Service action log. **Note:** On most systems with logical partitions, it is common to have one or more "Missing or Non-reporting" system bus resource's under Hardware Service Manager. See Hardware Service Manager for details.
There is a `Device Not Found` message during an installation from an alternate installation device.	Go to TUPIP06.

Table 7. Symptoms related to a logical partition on a server that has multiple logical partitions (continued)

Symptom	**Service action**
There is a problem with a guest partition. **Note:** These are problems reported from the operating system (other than IBM i) running in a guest partition or reported from the hosting partition of a guest partition.	If there are serviceable events in the logical partition or hosting partition, work on these problems first. If there are no SAL entries in the logical partition and no SAL entries in the hosting partition, contact your next level of support.

Table 8. Obvious physical symptoms

Symptom	**Service action**
One or more of the following was reported: • Noise • Smoke • Odor	Go to the system safety inspection procedures for your specific system.
A part is broken or damaged.	Go to the System parts to get the part number. Then go to the remove and replace procedures for your specific system to exchange the part.

Table 9. Time-of-day problems

Symptom	**Service action**
System clock loses or gains more than 1 second per day when the system is connected to utility power.	Replace the service processor. See symbolic FRU SVCPROC.
System clock loses or gains more than 1 second per day when the system is disconnected from utility power.	Replace the time-of-day battery on the service processor. Go to symbolic FRU TOD_BAT.

AIX server or AIX partition symptoms

Use the following tables to find the symptom you are experiencing. If you cannot find your symptom, contact your next level of support.

Choose the description that best describes your situation:

- You have a service action to complete
- An LED is not operating as expected
- Control (operator) panel problems
- Reference codes
- Management consoles
- There is a display or monitor problem (for example, distortion or blurring)
- Power and cooling problems
- Other symptoms or problems

You have a service action to complete

Symptom	What you should do:
You have an open service event in the service action event log.	Go to "Beginning problem analysis" on page 1.
You have parts to exchange or a corrective action to complete.	1. Go to the remove and replace procedures for your specific server. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	1. Go to Verifying the repair. 2. Go to the Close of call procedure.
You need to verify correct system operation.	1. Go to Verifying the repair. 2. Go to Close of call procedure.

An LED is not operating as expected

Symptom	What you should do:
The system attention LED on the control panel is on.	Go to "Beginning problem analysis" on page 1.
The rack indicator LED does not turn on, but a drawer identify LED is on.	1. Make sure that the rack indicator LED is properly mounted to the rack. 2. Make sure that the rack identify LED is properly cabled to the bus bar on the rack and to the drawer identify LED connector. 3. Replace the following parts one at a time: • Rack LED to bus bar cable • LED bus bar to drawer cable • LED bus bar 4. Contact your next level of support.

Control (operator) panel problems

Symptom	What you should do:
01 does not appear in the upper-left corner of the operator panel display after the power is connected and before you press the power-on button. Other symptoms appear in the operator panel display or LEDs before the power on button is pressed.	Go to Power problems.

Symptom	What you should do:
A bouncing or scrolling ball (moving row of dots) remains on the operator panel display, or the operator panel display is filled with dashes or blocks.	Verify that the operator panel connections to the system backplane are connected and properly seated. Also, reseat the service processor card. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Operator panel assembly. 2. Service processor.

Symptom	What you should do:
You have a blank display on the operator panel. Other LEDs on the operator panel appear to behave normally.	Verify that the operator panel connections to the system backplane are connected and properly seated. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. See Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Control (operator) panel assembly. 2. Service processor.
You have a blank display on the operator panel. Other LEDs on the operator panel are off.	Go to Power problems.

Reference codes

Symptom	What you should do:
An 8-digit error code is displayed.	Look up the reference code in the Reference codes section of IBM Knowledge Center. **Note:** If the repair for this code does not involve replacing a FRU (for instance, running an AIX command that fixes the problem or changing a hot-pluggable FRU), then update the AIX error log after the problem is resolved by completing the following steps: 1. In the online diagnostics, select **Task Selection** > **Log Repair Action**. 2. Select resource **sysplanar0**. On systems with a fault indicator LED, this changes the fault indicator LED from the fault state to the normal state.
The system stops with an 8-digit error code displayed when you boot.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that does **not** begin with 0 or 2.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that begins with 0 or 2.	Record SRN 101-*xxxx*, where *xxxx* is the 4-digit code that is displayed in the control panel. Then, look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.
The system stops and a 3-digit number is displayed on the control panel.	Add 101– to the left of the three digits to create an SRN, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN. If a location code is displayed under the 3-digit error code, look at the location to see whether it matches the failing component that the SRN pointed to. If they do not match, complete the action that is given in the error code table. If the problem still exists, replace the failing component from the location code. If a location code is displayed under the 3-digit error code, record the location code. Record SRN 101-*xxx*, where *xxx* is the 3-digit number that is displayed in the operator panel display, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.

Management consoles

Symptom	What you should do:
The management console cannot be used to manage a managed system, or the connection to the managed system is failing.	If the managed system is operating normally (that is, there are no error codes or other symptoms) the management console might have a problem, or the connection to the managed system might be damaged or incorrectly cabled. Complete the following steps: 1. Check the connections between the management console and the managed system. Correct any cabling errors, if found. If another cable is available, connect it in place of the existing cable and refresh the management console interface. You might have to wait up to 30 seconds for the managed system to reconnect. 2. Verify that any connected management console is connected to the managed system. **Note:** The managed system must have power that is connected, and either be waiting for a power-on instruction (that is, `01` is in the upper-left corner of the operator panel) or be running. If the managed system does not appear in the navigation area of the management console management environment, the management console or the connection to the managed system might be failing. 3. Go to the entry MAP: - Go to: Managing the HMC section. 4. If there is a problem with the service processor card or the system backplane, complete the following steps. - If you cannot fix the problem by using the HMC tests in the Managing the HMC section: a. Replace the service processor card. See the remove and replace procedures for your specific system. b. Replace the system backplane if not already replaced in substep a. See the remove and replace procedures for your specific system.

Symptom	What you should do:
The management console (HMC only) cannot call out by using the attached modem and the customer's telephone line.	If the managed system is operating normally (that is, there are no error codes or other symptoms), the management console might have a problem, or the connection to the modem and telephone line might have a problem. Complete the following steps: 1. Check the connections between the management console, the modem, and the telephone line. Correct any cabling errors, if found. 2. Go to the entry MAP in the Managing your server using the Hardware Management Console section.

There is a display problem (for example, distortion or blurring)

Symptom	What you should do:
All display problems.	1. If you are using the HMC: Go to the Managing the HMC section. 2. If you are using a graphics display, complete the following steps: a. Go to the problem determination procedures for the display. b. If you do not find a problem, complete the following steps, one at a time, until the problem is resolved: i) Replace the graphics display adapter. See the remove and replace procedures for your specific system. ii) Replace the backplane into which the card is plugged. See the remove and replace procedures for your specific system. 3. If you are using an ASCII terminal, complete the following steps: a. Make sure that the ASCII terminal is connected to S1. b. If problems persist, go to the problem determination procedures for the terminal. c. If you do not find a problem, replace the service processor. See the remove and replace procedures for your specific system.
There appears to be a display problem (distortion, blurring, and so on).	Go to the problem determination procedures for the display.

Power and cooling problems

Symptom	What you should do:
The system does not power on and no error codes are available.	Go to Power problems.
The power LEDs on the operator panel and the power supply do not come on or stay on.	1. Check the service processor error log. 2. Go to Power problems.
The power LEDs on the operator panel and the power supply come on and stay on, but the system does not power on.	1. Check the service processor error log. 2. Go to Power problems.
A rack or a rack-mounted unit will not power on.	1. Check the service processor error log. 2. Go to Power problems.
The cooling fans do not come on, or come on but do not stay on.	1. Check the service processor error log. 2. Go to Power problems.
The system attention LED on the operator panel is on and there is no error code displayed.	1. Check the service processor error log. 2. Go to Power problems.

Other symptoms or problems

Symptom	What you should do:
The system stopped and a code is displayed on the operator panel.	Go to “Beginning problem analysis” on page 1.
`01` is displayed in the upper-left corner of the operator panel and the fans are off.	The service processor is ready. The system is waiting for power-on. Boot the system. If the boot is unsuccessful, and the system returns to the default display (indicated by `01` in the upper-left corner of the operator panel), go to MAP 0200: Problem determination procedure.
The operator panel displays STBY.	The service processor is ready. The server was shut down by the operating system and is still powered on. This condition can be requested by a privileged system user with no faults. Go to “Beginning problem analysis” on page 1. **Note:** See the service processor error log for possible operating system fault indications.
All of the system power-on self-test (POST) indicators are displayed on the firmware console, the system pauses and then restarts. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the POST.	Go to Problems with loading and starting the operating system.

Symptom	**What you should do:**
The system stops and all of the POST indicators are displayed on the firmware console. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the power-on self-test (POST).	Go to Problems with loading and starting the operating system.
The system stops and the message `starting software please wait...`is displayed on the firmware console.	Go to Problems with loading and starting the operating system.
The system does not respond to the password that you entered or the system login prompt is displayed when you boot in service mode.	1. If the password is being entered from HMC: Go to the Managing the HMC. 2. If the password is being entered from a keyboard that is attached to the system, the keyboard or its controller might be faulty. In this case, replace these parts in the following order: a. Keyboard b. Service processor 3. If the password is being entered from an ASCII terminal, use the problem determination procedures for the ASCII terminal. Make sure that the ASCII terminal is connected to S1. If the problem persists, replace the service processor. If the problem is fixed, go to “MAP 0410: Repair checkout” on page 32.
The system stops with a prompt to enter a password.	Enter the password. You cannot continue until a correct password is entered. After you enter a valid password, go to the beginning of this table and wait for one of the other conditions to occur.
The system does not respond when the password is entered.	Go to Step 1020-2.
No codes are displayed on the operator panel within a few seconds of turning on the system. The operator panel is blank before the system is powered on.	Reseat the operator panel cable. If the problem is not resolved, replace in the following order: 1. Operator panel assembly. See the remove and replace procedures for your specific system. 2. Service processor. See the remove and replace procedures for your specific system. If the problem is fixed, go to “MAP 0410: Repair checkout” on page 32. If the problem is still not corrected, go to MAP 0200: Problem determination procedure.

Symptom	**What you should do:**
The SMS configuration list or boot sequence selection menu shows more SCSI devices that are attached to a controller or an adapter than are actually attached.	A device might be set to use the same SCSI bus ID as the control adapter. Note the ID being used by the controller or adapter (this can be checked or changed through an SMS utility), and verify that no device that is attached to the controller is set to use that ID. If settings do not appear to be in conflict, complete the following steps: 1. Go to MAP 0200: Problem determination procedure. 2. Replace the SCSI cable. 3. Replace the device. 4. Replace the SCSI adapter **Note:** In a **twin-tailed** configuration where there is more than one initiator device (normally another system) attached to the SCSI bus, it might be necessary to use SMS utilities to change the ID of the SCSI controller or adapter.
You have a problem that does not prevent the system from booting. The operator panel is functional and the rack indicator LED operates as expected.	Go to MAP 0200: Problem determination procedure.
All other symptoms.	Go to MAP 0200: Problem determination procedure.
All other problems.	Go to MAP 0200: Problem determination procedure.
You do not have a symptom.	Go to MAP 0200: Problem determination procedure.
You have parts to exchange or a corrective action to complete.	1. Go to “Beginning problem analysis” on page 1. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	Go to “MAP 0410: Repair checkout” on page 32.
You need to verify correct system operation.	Go to “MAP 0410: Repair checkout” on page 32.

Symptom	**What you should do:**
The system stopped. A POST indicator is displayed on the system console and an eight-digit error code is not displayed.	If the POST indicator represents: 1. Memory, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 2. Keyboard a. Replace the keyboard. b. Replace the service processor, location: model dependent. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure. 3. Network, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 4. SCSI, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 5. Speaker a. Replace the control panel. The location depends on the model. b. Replace the service processor. The location depends on the model. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The diagnostic operating instructions are displayed.	Go to MAP 0200: Problem determination procedure.
The system login prompt is displayed.	If you are loading the diagnostics from a CD-ROM, you might not have pressed the correct key or you might not have pressed the key soon enough when you were trying to indicate a service mode IPL of the diagnostic programs. If so, try to boot the CD-ROM again and press the correct key. **Note:** Complete the system shutdown procedure before you turn off the system. If you are sure that you pressed the correct key in a timely manner, go to Step 1020-2. If you are loading diagnostics from a Network Installation Management (NIM) server, check for the following items: • The bootlist on the client might be incorrect. • Cstate on the NIM server might be incorrect. • Network problems might be preventing you from connecting to the NIM server. Verify the settings and the status of the network. If you continue to have problems see Problems with loading and starting the operating system and follow the steps for network boot problems.

Symptom	What you should do:
The System Management Services (SMS) menu is displayed when you were trying to boot stand-alone AIX diagnostics.	If you are loading diagnostics from the CD-ROM, you might not have pressed the correct key when you were trying to indicate a service mode IPL of the diagnostic programs. If so, try to boot the CD-ROM again and press the correct key. If you are sure that you pressed the correct key, the device or media you are attempting to boot from might be faulty. 1. Try to boot from an alternate boot device that is connected to the same controller as the original boot device. If the boot succeeds, replace the original boot device (for removable media devices, try the media first). If the boot fails, go to Problems with loading and starting the operating system. 2. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The SMS boot sequence selection menu or remote IPL menu does not show all of the bootable devices in the partition or system.	If an AIX or Linux partition is being booted, verify that the devices that you expect to see in the list are assigned to this partition. If they are not, use the management console to reassign the required resources. If they are assigned to this partition, go to Problems with loading and starting the operating system to resolve the problem.

MAP 0410: Repair checkout

Use this MAP to check out the server after a repair is completed.

Purpose of this MAP

Use this MAP to check out the server after a repair is completed.

Note: Only use standalone diagnostics for repair verification when no other diagnostics are available on the system. Standalone diagnostics do not log repair actions.

If you are servicing an SP system, go to the End-of-call MAP in the *SP System Service Guide*.

If you are servicing a clustered system, go to the End of Call MAP in the *Clustered eServer Installation and Service Guide*.

If you are servicing a cluster, go to the End-of-call procedure in the *Clustered Installation and Service Guide*.

- **Step 0410-1**

  Did you use an AIX diagnostics service aid hot-swap operation to change the FRU?

  No

  Go to Step 0410-2.

  Yes

  Go to Step 0410-4.

- **Step 0410-2**

  Note: If the system backplane or battery has been replaced and you are loading diagnostics from a server over a network, it may be necessary for the customer to set the network boot information for this system before diagnostics can be loaded. The system time and date information should also be set when the repair is completed.

Do you have any FRUs (for example cards, adapters, cables, or devices) that were removed during problem analysis that you want to put back into the system?

No
Go to Step 0410-3.

Yes
Reinstall all of the FRUs that were removed during problem analysis. Go to Step 0410-3.

- **Step 0410-3**

  Is the system or logical partition that you are performing a repair action on running the AIX operating system?

  No
  Go to Step 0410-5.

  Yes
  Go to Step 0410-4.

- **Step 0410-4**

  Does the system or logical partition you are performing a repair action on have AIX installed?

  Note: Answer **No** to this question if you have just replaced a hard disk in the root volume group.

  No
  Go to Step 0410-5.

  Yes
  Go to Step 0410-6.

- **Step 0410-5**

  Run standalone diagnostics from either a CD ROM or from a NIM server.

  Did you encounter any problems?

  No
  Go to Step 0410-14.

  Yes
  Go to MAP 0020: Problem determination procedure.

- **Step 0410-6**

  1. Power on the system.
  2. Wait until the AIX operating system login prompt displays or until system activity on the operator panel or display apparently has stopped.

  Did the AIX Login Prompt display?

  No
  Go to MAP 0020: Problem determination procedure.

  Yes
  Go to Step 0410-7.

- **Step 0410-7**

  If the **Resource Repair Action** menu is already displayed, go to Step 0410-10; otherwise, complete the following steps:

  1. Log into the operating system either with root authority (if needed, ask the customer to enter the password) or use the CE login.
  2. Enter the `diag -a` command and check for missing resources. Follow any instructions that display. If an SRN displays, suspect a loose card or connection. If no instructions display, no resources were detected as missing. Continue on to Step 0410-8

- **Step 0410-8**

  1. Enter `diag` at the command prompt.
  2. Press Enter.
  3. Select the **Diagnostics Routines** option.
  4. When the DIAGNOSTIC MODE SELECTION menu displays, select **Problem determination**.
  5. When the ADVANCED DIAGNOSTIC SELECTION menu displays, select the **All Resources** option or test the FRUs you exchanged, and any devices that are attached to the FRUs you exchanged, by selecting the diagnostics for the individual FRU.

  Did the RESOURCE REPAIR ACTION menu (801015) display?

  No
  Go to Step 0410-9.

  Yes
  Go to Step 0410-10.

- **Step 0410-9**

  Did the TESTING COMPLETE, no trouble was found menu (801010) display?

  No
  There is still a problem. Go to MAP 0020: Problem determination procedure.

  Yes
  Use the **Log Repair Action** option, if not previously logged, in the TASK SELECTION menu to update the AIX error log. If the repair action was reseating a cable or adapter, select the resource associated with that repair action.

  If the resource associated with your action is not displayed on the resource list, select **sysplanar0**.

  Note: If the check log indicator is on, this will set it back to the normal state.

  Go to Step 0410-12.

- **Step 0410-10**

  When a test is run on a resource in system verification mode, and that resource has an entry in the AIX error log, if the test on the resource was successful, the RESOURCE REPAIR ACTION menu displays.

  After replacing a FRU, you must select the resource for that FRU from the RESOURCE REPAIR ACTION menu. This updates the AIX error log to indicate that a system-detectable FRU has been replaced.

  Note: If the check log indicator is on, this action will set it back to the normal state.

  Complete the following steps:

  1. Select the resource that has been replaced from the RESOURCE REPAIR ACTION menu. If the repair action was reseating a cable or adapter, select the resource associated with that repair action. If the resource associated with your action is not displayed on the resource list, select **sysplanar0**.
  2. Press **Commit** after you make your selections.

  Did another Resource Repair Action (801015) display?

  No
  If the No Trouble Found menu displays, go to Step 0410-12.

  Yes
  Go to Step 0410-11.

- **Step 0410-11**

  The parent or child of the resource you just replaced may also require that you run the RESOURCE REPAIR ACTION service aid on it.

  When a test is run on a resource in system verification mode, and that resource has an entry in the AIX error log, if the test on the resource was successful, the RESOURCE REPAIR ACTION menu displays.

After replacing that FRU, you must select the resource for that FRU from the RESOURCE REPAIR ACTION menu. This updates the AIX error log to indicate that a system-detectable FRU has been replaced.

Note: If the check log indicator is on, this action will set it back to the normal state.

Complete the following steps:

1. From the RESOURCE REPAIR ACTION menu, select the parent or child of the resource that has been replaced . If the repair action was reseating a cable or adapter, select the resource associated with that repair action. If the resource associated with your action is not displayed on the resource list, select **sysplanar0**.
2. Press COMMIT after you make your selections.
3. If the No Trouble Found menu displays, go to Step 0410-12.

- **Step 0410-12**

  If you changed the service processor or network settings, as instructed in previous MAPs, restore the settings to the value they had prior to servicing the system. If you ran standalone diagnostics from CD-ROM, remove the standalone diagnostics CD-ROM from the system.

  Did you perform service on a RAID subsystem involving changing of the PCI RAID adapter cache card or changing the configuration?

  No
  Go to Step 0410-14.

  Yes
  Go to Step 0410-13.

- **Step 0410-13**

  Use the **Recover Options** selection to resolve the RAID configuration. To do this, complete the following steps:

  1. On the PCI SCSI Disk Array Manager screen, select **Recovery options**.
  2. If a previous configuration exists on the replacement adapter, this must be cleared. Select **Clear PCI SCSI Adapter Configuration**. Press F3.
  3. On the Recovery Options screen, select **Resolve PCI SCSI RAID Adapter Configuration**.
  4. On the Resolve PCI SCSI RAID Adapter Configuration screen, select **Accept Configuration on Drives**.
  5. On the PCI SCSI RAID Adapter selections menu, select the adapter that you changed.
  6. On the next screen, press Enter.
  7. When you get the Are You Sure selection menu, press Enter to continue.
  8. You should get an OK status message when the recover is complete. If you get a `Failed` status message, verify that you selected the correct adapter, then repeat this procedure. When recover is complete, exit the operating system.
  9. Go to Step 0410-14.

- **Step 0410-14**

  If the system you are servicing has a management console, go to the end-of-call procedure for systems with Service Focal Point.

This completes the repair; return the server to the customer.

Linux server or Linux partition symptoms

Use the following tables to find the symptom you are experiencing. If you cannot find your symptom, contact your next level of support.

Choose the description that best describes your situation:

- You have a service action to complete
- An LED is not operating as expected
- Control (operator) panel problems
- Reference codes
- Management consoles
- There is a display or monitor problem (for example, distortion or blurring)
- Power and cooling problems
- Other symptoms or problems

You have a service action to complete

Symptom	What you should do:
You have an open service event in the service action event log.	Go to "Beginning problem analysis" on page 1.
You have parts to exchange or a corrective action to complete.	1. Go to the remove and replace procedures for your specific system. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	1. Go to Verifying the repair. 2. Go to the Close of call procedure.
You need to verify correct system operation.	1. Go to Verifying the repair. 2. Go to the Close of call procedure.

An LED is not operating as expected

Symptom	What you should do:
The system attention LED on the control panel is on.	Go to "Linux fast-path problem isolation" on page 46.
The rack identify LED does not operate properly.	Go to the "Linux fast-path problem isolation" on page 46.
The rack indicator LED does not turn on, but a drawer identify LED is on.	1. Make sure that the rack indicator LED is properly mounted to the rack. 2. Make sure that the rack identify LED is properly cabled to the bus bar on the rack and to the drawer identify LED connector. 3. Replace the following parts one at a time: • Rack LED to bus bar cable • LED bus bar to drawer cable • LED bus bar 4. Contact your next level of support.

Control (operator) panel problems

Symptom	What you should do:
01 does not appear in the upper-left corner of the operator panel display after the power is connected and before you press the power-on button. Other symptoms appear in the operator panel display or LEDs before the power on button is pressed.	Go to Power problems.
A bouncing or scrolling ball (moving row of dots) remains on the operator panel display, or the operator panel display is filled with dashes or blocks.	Verify that the operator panel connections to the system backplane are connected and properly seated. Also, reseat the service processor card. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Operator panel assembly. 2. Service processor.

Symptom	**What you should do:**
You have a blank display on the operator panel. Other LEDs on the operator panel appear to behave normally.	Verify that the operator panel connections to the system backplane are connected and properly seated. If a client computer (such as a PC with Ethernet capability and a Web browser) is available, connect it to the service processor in the server that is displaying the symptom. To connect a personal computer with Ethernet capability and a Web browser, or an ASCII terminal, to access the Advanced System Management Interface (ASMI), go to Managing your server using the Advanced System Management Interface. • If you can successfully access the ASMI, replace the operator panel assembly. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. • If you cannot successfully access the ASMI, replace the service processor. Refer to Finding parts, locations, and addresses to determine the part number and correct exchange procedure. If you do not have a PC or ASCII terminal, replace the following one at a time (go to Finding parts, locations, and addresses to determine the part number and correct exchange procedure): 1. Control (operator) panel assembly. 2. Service processor.
You have a blank display on the operator panel. Other LEDs on the operator panel are off.	Go to Power problems.

Reference codes

Symptom	What you should do:
An 8-digit error code is displayed.	Look up the reference code in the Reference codes section of IBM Knowledge Center. **Note:** If the repair for this code does not involve replacing a FRU (for instance, running an AIX command that fixes the problem or changing a hot-pluggable FRU), then update the AIX error log after the problem is resolved by completing the following steps: 1. In the online diagnostics, select **Task SelectionLog Repair Action**. 2. Select resource **sysplanar0**. On systems with a fault indicator LED, this changes the "fault indicator" LED from the "fault" state to the "normal" state.
The system stops with an 8-digit error code displayed when you boot.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that does **not** begin with 0 or 2.	Look up the reference code in the Reference codes section of IBM Knowledge Center.
The system stops and a 4-digit code displays on the control panel that begins with 0 or 2 is displayed in the operator panel display.	Record SRN 101-*xxxx*, where *xxxx* is the 4-digit code that is displayed in the control panel, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.
The system stops and a 3-digit number is displayed on the control panel.	Add 101– to the left of the three digits to create an SRN, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN. If a location code is displayed under the 3-digit error code, look at the location to see whether it matches the failing component that the SRN pointed to. If they do not match, complete the action that is given in the error code table. If the problem still exists, then replace the failing component from the location code. If a location code is displayed under the 3-digit error code, record the location code. Record SRN 101-*xxx*, where *xxx* is the 3-digit number that is displayed in the operator panel display, then look up this reference code in the Reference codes section of IBM Knowledge Center. Follow the instructions that are given in the Description and Action column for your SRN.

Management consoles

Symptom	**What you should do:**
The management console cannot be used to manage a managed system, or the connection to the managed system is failing.	If the managed system is operating normally (no error codes or other symptoms), the management console might have a problem, or the connection to the managed system might be damaged or incorrectly cabled. Complete the following steps: 1. Check the connections between the management console and the managed system. Correct any cabling errors if found. If another cable is available, connect it in place of the existing cables and refresh the management console interface. You might need to wait up to 30 seconds for the managed system to reconnect. 2. Verify that any connected management console is connected to the managed system by checking the Management Environment of the management console. **Note:** The managed system must have power that is connected and the system running, or waiting for a power-on instruction (`01` is in the upper-left corner of the operator panel). If the managed system does not appear in the Navigation area of the management console Management Environment, the management console or the connection to the managed system might be failing. 3. Go to the Managing the HMC section. 4. There might be a problem with the service processor card or the system backplane. If you cannot fix the problem by using the HMC tests in the Managing the HMC section, complete the following steps, one at a time, until the problem is resolved: a. Replace the service processor card. Refer to the remove and replace procedures for your specific system. b. Replace the management console system backplane. Refer to the remove and replace procedures for your specific system.

Symptom	What you should do:
The management console (HMC only) cannot call out using the attached modem and the customer's telephone line.	If the managed system is operating normally (no error codes or other symptoms), the management console might have a problem, or the connection to the modem and telephone line might have a problem. Complete the following steps: 1. Check the connections between the management console, the modem, and the telephone line. Correct any cabling errors, if found. 2. Go to the Managing your server using the Hardware Management Console section.

There is a display problem (for example, distortion or blurring)

Symptom	What you should do:
All display problems.	1. Go to the Managing the HMC section. 2. If you are using a graphics display, complete the following steps: a. Go to the problem determination procedures for the display. b. If you do not find a problem, complete the following steps, one at a time, until the problem is resolved: i) Replace the graphics display adapter. Refer to the remove and replace procedures for your specific system. ii) Replace the backplane into which the graphics display adapter is plugged. Refer to the remove and replace procedures for your specific system.
There appears to be a display problem (distortion, blurring, and so on).	Go to the problem determination procedures for the display.

Power and cooling problems

Symptom	What you should do:
The system does not power on and no error codes are available.	Go to Power problems.
The power LEDs on the operator panel and the power supply do not come on or stay on.	1. Check the service processor error log. 2. Go to Power problems.
The power LEDs on the operator panel and the power supply come on and stay on, but the system does not power on.	1. Check the service processor error log. 2. Go to Power problems.
A rack or a rack-mounted unit will not power on.	1. Check the service processor error log. 2. Go to Power problems.

Symptom	What you should do:
The cooling fans do not come on, or come on but do not stay on.	1. Check the service processor error log. 2. Go to Power problems.
The system attention LED on the operator panel is on and no error code is displayed.	1. Check the service processor error log. 2. Go to Power problems.

Other symptoms or problems

Symptom	What you should do:
The system stopped and a code is displayed on the operator panel.	Go to "Beginning problem analysis" on page 1.
`01` is displayed in the upper-left corner of the operator panel and the fans are off.	The service processor is ready. The system is waiting for power-on. Boot the system. If the boot is unsuccessful, and the system returns to the default display (indicated by `01` in the upper-left corner of the operator panel), go to MAP 0020: Problem determination procedure.
The operator panel displays STBY.	The service processor is ready. The server was shut down by the operating system and is still powered on. This condition can be requested by a privileged system user with no faults. Go to "Beginning problem analysis" on page 1. **Note:** See the service processor error log for possible operating system fault indications.
All of the system POST indicators are displayed on the firmware console, the system pauses and then restarts. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the power-on self-test (POST).	Go to Problems with loading and starting the operating system.
The system stops and all of the POST indicators are displayed on the firmware console. The term *POST indicators* refers to the device mnemonics (the words `memory`, `keyboard`, `network`, `scsi`, and `speaker`) that appear on the firmware console during the power-on self-test (POST).	Go to Problems with loading and starting the operating system.
The system stops and the message `starting software please wait...`is displayed on the firmware console.	Go to Problems with loading and starting the operating system.

Symptom	What you should do:
The system does not respond to the password that was entered or the system login prompt is displayed when you boot the system in service mode.	1. Go to the Managing the HMC. 2. If the password is being entered from a keyboard that is attached to the system, the keyboard or its controller might be faulty. In this case, replace these parts in the following order: a. Keyboard b. Service processor If the problem is fixed, go to “MAP 0410: Repair checkout” on page 32.
The system stops with a prompt to enter a password.	Enter the password. You cannot continue until a correct password is entered. After you enter a valid password, go to the beginning of this table and wait for one of the other conditions to occur.
The system does not respond when the password is entered.	Go to Step 1020-2.
No codes are displayed on the operator panel within a few seconds of turning on the system. The operator panel is blank before the system is powered on.	Reseat the operator panel cable. If the problem is not resolved, replace in the following order: 1. Operator panel assembly. Refer to the remove and replace procedures for your specific system. 2. Service processor. Refer to the remove and replace procedures for your specific system. If the problem is fixed, go to “MAP 0410: Repair checkout” on page 32. If the problem is still not corrected, go to MAP 0020: Problem determination procedure.
The SMS configuration list or boot sequence selection menu shows more SCSI devices that are attached to a controller/adapter than are actually attached.	A device might be set to use the same SCSI bus ID as the control adapter. Note the ID being used by the controller/adapter (this can be checked or changed through an SMS utility), and verify that no device that is attached to the controller is set to use that ID. If settings do not appear to be in conflict, complete the following steps: 1. Go to MAP 0020: Problem determination procedure. 2. Replace the SCSI cable. 3. Replace the device. 4. Replace the SCSI adapter. **Note:** In a "twin-tailed" configuration where there is more than one initiator device (normally another system) attached to the SCSI bus, it might be necessary to use SMS utilities to change the ID of the SCSI controller or adapter.

Symptom	What you should do:
You suspect a cable problem.	Go to Adapters, Devices and Cables for Multiple Bus Systems.
You have a problem that does not prevent the system from booting. The operator panel is functional and the rack indicator LED operates as expected.	Go to MAP 0020: Problem determination procedure.
All other symptoms.	Go to MAP 0020: Problem determination procedure.
All other problems.	Go to MAP 0020: Problem determination procedure.
You do not have a symptom.	Go to MAP 0020: Problem determination procedure.
You have parts to exchange or a corrective action to complete.	1. Go to "Beginning problem analysis" on page 1. 2. Go to Close of call procedure.
You need to verify that a part exchange or corrective action corrected the problem.	Go to "MAP 0410: Repair checkout" on page 32.
You need to verify correct system operation.	Go to "MAP 0410: Repair checkout" on page 32.
The system stopped. A POST indicator is displayed on the system console and an eight-digit error code is not displayed.	If the POST indicator represents: 1. Memory, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 2. Keyboard a. Replace the keyboard. b. Replace the service processor. The location depends on the model. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure. 3. Network, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 4. SCSI, go to PFW 1548: Memory and processor subsystem problem isolation procedure. 5. Speaker a. Replace the control panel. The location depends on the model. b. Replace the service processor. The location depends on the model. c. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The diagnostic operating instructions are displayed.	Go to MAP 0020: Problem determination procedure.

Symptom	What you should do:
The system login prompt is displayed.	If you are loading the diagnostics from a CD-ROM, you might not have pressed the correct key or you might not have pressed the key soon enough when you were trying to indicate a service mode IPL of the diagnostic programs. If so, start again at the beginning of this step. **Note:** Complete the system shutdown procedure before you turn off the system. If you are sure that you pressed the correct key in a timely manner, go to Step 1020-2. If you are loading diagnostics from a Network Installation Management (NIM) server, check for the following items: • The bootlist on the client might be incorrect. • Cstate on the NIM server might be incorrect. • There might be network problems that are preventing you from connecting to the NIM server. Verify the settings and the status of the network. If you continue to have problems refer to Problems with loading and starting the operating system and follow the steps for network boot problems.
The System Management Services (SMS) menu is displayed when you were trying to boot stand-alone diagnostics.	If you are loading diagnostics from the CD-ROM, you might not have pressed the correct key when you were trying to indicate a service mode IPL of the diagnostic programs. If so, try to boot the CD-ROM again and press the correct key. If you are sure that you pressed the correct key, the device or media you are attempting to boot from might be faulty. 1. Try to boot from an alternate boot device that is connected to the same controller as the original boot device. If the boot succeeds, replace the original boot device (for removable media devices, try the media first). If the boot fails, go to problems with loading and starting the operating system. 2. Go to PFW 1548: Memory and processor subsystem problem isolation procedure.
The SMS boot sequence selection menu or remote IPL menu does not show all of the bootable devices in the partition or system.	If an AIX or Linux partition is being booted, verify that the devices that you expect to see in the list are assigned to this partition. If they are not, use the management console to reassign the required resources. If they are assigned to this partition, go to Problems with loading and starting the operating system to resolve the problem.

Linux fast-path problem isolation

Use this information to help you isolate a hardware problem when you use the Linux operating system.

Linux fast path table

Locate the problem in the following table and then go to the action indicated for the problem.

Symptoms	Action
You have an eight-digit reference code.	Go to Reference codes, and do the listed action for the eight-digit reference code.
You are trying to isolate a problem on a Linux server or a partition that is running Linux operating system.	**Note:** This procedure is used to help display an eight-digit reference code by using system log information. Before you use this procedure, if you are having a problem with a media device such as a tape or DVD-ROM drive, continue through this table and follow the actions for the appropriate device. Go to "Linux problem isolation procedure" on page 49.
You suspect a problem with your server but you do not have any specific symptom.	Go to MAP 0020: Problem determination procedure for problem determination procedures.
You need to run the eServer™ stand-alone diagnostics.	Go to AIX and Linux diagnostics and service aids
SRNs	
You have an SRN.	Look up the SRN in the Service request numbers and do the listed action.
An SRN is displayed when you run the eServer stand-alone diagnostics.	1. Record the SRN and location code. 2. Look up the SRN in the Service request numbers and do the listed action.
Tape Drive Problems	
You suspect a tape drive problem.	1. Refer to the tape drive documentation and clean the tape drive. 2. Refer to the tape drive documentation and do any listed problem determination procedures. 3. Go to MAP 0020: Problem determination procedure for problem determination procedures. **Note:** Information on tape cleaning and tape-problem determination is normally either in the tape drive operator guide or the system operator guide.
Optical Drive Problems	
You suspect an optical drive problem.	1. Refer to the optical documentation and do any listed problem determination procedures. 2. Before you service an optical drive, ensure that it is not in use and that the power connector is correctly attached to the drive. If the load or unload operation does not function, replace the optical drive. 3. Go to MAP 0020: Problem determination procedure for problem determination procedures. **Note:** If the optical drive has its own user documentation, follow any problem determination for the optical drive.

Symptoms	Action
SCSI Disk Drive Problems	
You suspect a disk drive problem. Disk problems are logged in the error log and are analyzed when the stand-alone disk diagnostics are run in problem determination mode. Problems are reported if the number of errors is above defined thresholds.	Go to MAP 0020: Problem determination procedure for problem determination procedures.
Token-Ring Problems	
You suspect a token-ring adapter or network problem.	1. Check with the network administrator for known problems. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures.
Ethernet Problems	
You suspect an Ethernet adapter or network problem.	1. Check with the network administrator for known problems. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures.
Display Problems	
You suspect a display problem.	1. If your display is connected to a KVM switch, go to Troubleshooting the keyboard, video, and mouse (KVM) switch for the 1x8 and 2x8 console manager. If you are still having display problems after you complete the KVM switch procedures, come back here and continue with step 2. 2. Go to the Managing your server using the Hardware Management Console section. 3. If you are using a graphics display, complete the following steps: a. Go to the problem determination procedures for the display. b. If you do not find a problem, complete the following steps, one at a time, until the problem is resolved: i) Replace the graphics display adapter. Refer to the remove and replace procedures for your specific system. ii) Replace the backplane into which the graphics display adapter is plugged. Refer to the remove and replace procedures for your specific system.
Keyboard or Mouse	

Symptoms	Action
You suspect a keyboard or mouse problem.	If your keyboard is connected to a KVM switch, go to Troubleshooting the keyboard, video, and mouse (KVM) switch for the 1x8 and 2x8 console manager. If you are still having keyboard problems after you complete the KVM switch procedures, come back here and continue to the next paragraph. Go to MAP 0020: Problem determination procedure for problem determination procedures. If you are unable to run diagnostics because the system does not respond to the keyboard, replace the keyboard or system backplane. **Note:** If the problem is with the keyboard, it might be caused by the mouse device. To check, unplug the mouse and then recheck the keyboard. If the keyboard works, replace the mouse.
System Messages	
A System Message is displayed.	1. If the message describes the cause of the problem, attempt to correct it. 2. Look for another symptom to use.
System Hangs or Loops when you run the OS or Diagnostics	
The system hangs in the same application.	Suspect the application. To check the system, complete the following steps: 1. Power off the system. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures. 3. If an SRN is displayed at anytime, record the SRN and location code. 4. Look up the SRN in the Service request numbers and do the listed action.
The system hangs in various applications.	1. Power off the system. 2. Go to MAP 0020: Problem determination procedure for problem determination procedures. 3. If an SRN is displayed at anytime, record the SRN and location code. 4. Look up the SRN in the Service request numbers and do the listed action.
The system hangs when you run diagnostics.	Replace the resource that is being tested.
Exchanged FRUs Did Not Fix the Problem	
A FRU or FRUs you exchanged did not fix the problem.	Go to MAP 0020: Problem determination procedure for problem determination procedures.
You Cannot Find the Symptom in This Table	
All other problems.	Go to MAP 0020: Problem determination procedure for problem determination procedures.

Linux problem isolation procedure
Use this procedure when servicing a Linux partition or a server that has Linux as its only operating system.

About this task

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

These procedures define the steps to take when servicing a Linux partition or a server that has Linux as its only operating system.

Before continuing with this procedure it is recommended that you review the additional software available to enhance your Linux solutions. See Service and productivity tools for PowerLinux servers (http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags).

Note: If the server is attached to a management console, the various codes that might display on the management console are all listed as reference codes by Service Focal Point (SFP). Use the following table to help you identify the type of error information that might be displayed when you are using this procedure.

Number of digits in reference code	Reference code	Name or code type
Any	Contains # (number sign)	Menu goal
Any	Contains - (hyphen)	Service request number (SRN)
5	Does not contain # or -	SRN
8	Does not contain # or -	system reference code (SRC)

Procedure

1. Is the server managed by a management console that is running Service Focal Point (SFP)?

   No
   Go to step "3" on page 50.

   Yes
   Go to step "2" on page 50.
2. Servers with Service Focal Point

   Look at the service action event log in SFP for errors. Focus on those errors with a timestamp near the time at which the error occurred. Follow the steps indicated in the error log entry to resolve the problem. If the problem is not resolved, continue with step "3" on page 50.
3. Look for and record all reference code information or software messages on the operator panel and in the service processor error log (which is accessible by viewing the ASMI menus).
4. Choose a Linux partition that is running correctly (preferably the partition with the problem).

   Is Linux usable in any partition with Linux installed?

   No
   Go to step "10" on page 51.

   Yes
   Go to step "5" on page 50.
5. Diagnose the RTAS events. For instructions, see Diagnosing RTAS events.
6. Record any RTAS events found in the Linux system log

   If the system is configured with more than one logical partition with Linux installed, repeat step "5" on page 50 and step "6" on page 50 for all logical partitions that have Linux installed.
7. Examine the Linux boot (IPL) log by logging in to the system as the root user and entering the following command:

   ```
   cat /var/log/boot.msg |grep RTAS |more
   ```

   Linux boot (IPL) error messages are logged into the **boot.msg** file under **/var/log**. An example of the Linux boot error log:

   ```
   RTAS daemon started
   RTAS: -------- event-scan begin --------
   RTAS: Location Code: U0.1-F3
   RTAS: WARNING: (FULLY RECOVERED) type: SENSOR
   RTAS: initiator: UNKNOWN target: UNKNOWN
   RTAS: Status: bypassed new
   RTAS: Date/Time: 20020830 14404000
   RTAS: Environment and Power Warning
   RTAS: EPOW Sensor Value: 0x00000001
   RTAS: EPOW caused by fan failure
   RTAS: -------- event-scan end ----------
   ```

8. Record any RTAS events found in the Linux boot (IPL) log in step “7” on page 50.

   Ignore all other events in the Linux boot (IPL) log. If the system is configured with more than one logical partition with Linux installed, repeat step “7” on page 50 and step “8” on page 51 for all logical partitions that have Linux installed.
9. Record any extended data found in the Linux system log in Step “5” on page 50 or the Linux boot (IPL) log in step “7” on page 50.

   Note: The lines in the Linux extended data that begin with `<4>RTAS: Log Debug: 04` contain the reference code listed in the next 8 hexadecimal characters. In the previous example, `4b27 26fb` is a reference code. The reference code is also known as word 11. Each 4 bytes after the reference code in the Linux extended data is another word (for example, `04a0 0011` is word 12, and `702c 0014` is word 13, and so on).

   If the system is configured with more than one logical partition with Linux installed, repeat step “9” on page 51 for all logical partitions that have Linux installed.
10. Were any reference codes or checkpoints recorded in steps “3” on page 50, “6” on page 50, “8” on page 51, or “9” on page 51?

    No

    Go to step “11” on page 51.

    Yes

    Go to the Linux fast-path problem isolation with each reference code that was recorded. Perform the indicated actions one at a time for each reference code until the problem has been corrected. If all recorded reference codes have been processed and the problem has not been corrected, go to step “11” on page 51.
11. If no additional error information is available and the problem has not been corrected, complete the following steps:

    a) Shut down the system.

    b) If a management console is not attached, see Managing your server using the Advanced System Management Interface for instructions to access the ASMI.

    Note: The ASMI functions can also be accessed by using a personal computer connected to system port 1.

    You need a personal computer capable of connecting to system port 1 on the system unit. (The Linux login prompt cannot be seen on a personal computer connected to system port 1.) If the ASMI functions are not otherwise available, use the following procedure:

    i) Attach the personal computer and cable to system port 1 on the system unit.

    ii) With 01 displayed in the operator panel, press a key on the virtual terminal on the personal computer. The service ASMI menus are available on the attached personal computer.

    iii) If the service processor menus are not available on the personal computer, perform the following steps:

    a) Examine and correct all connections to the service processor.

    b) Replace the service processor.

    Note: The service processor might be contained on a separate card or board; in some systems, the service processor is built into the system backplane. Contact your next level of support for help before replacing a system backplane.

    c) Examine the service processor error log.

    Record all reference codes and messages written to the service processor error log. Go to step “12” on page 51.
12. Were any reference codes recorded in step “11” on page 51?

    No

    Go to step “20” on page 53.

Yes

Go to the Linux fast-path problem isolation with each reference code or symptom you have recorded. Perform the indicated actions, one at a time, until the problem has been corrected. If all recorded reference codes have been processed and the problem has not been corrected, go to "20" on page 53.

13. Reboot the system and bring all partitions to the login prompt.

    If Linux is not usable in all partitions, go to step "17" on page 53.

14. Use the `lscfg` command to list all resources assigned to all partitions.

    Record the adapter and the partition for each resource.

15. To determine whether any devices or adapters are missing, compare the list of partition assignments, and resources found, to the customer's known configuration. Record the location of any missing devices.

    Also record any differences in the descriptions or the locations of devices.

    You may also compare this list of resources that were found to an earlier version of the device tree as follows:

    Note: At the Linux command prompt, type `vpdupdate`, and press Enter. The device tree is stored in the `/var/lib/lsvpd/` directory in a file with the file name device-tree-YYYY-MM-DD-HH:MM:SS, where YYYY is the year, MM is the month, DD is the day, and HH, MM, and SS are the hour, minute and second, respectively, of the date of creation.

    - At the command line, type the following:

      ```
      cd /var/lib/lsvpd/
      ```

    - At the command line, type the following:

      ```
      lscfg -vpz /var/lib/lsvpd/<file_name>
      ```

      Where, *<file_name>* is the .gz file name that contains the database archive.

    The **diff** command offers a way to compare the output from a current **lscfg** command to the output from an older **lscfg** command. If the files names for the current and old device trees are **current.out** and **old.out**, respectively, type: `diff old.out current.out`. Any lines that exist in the old, but not in the current will be listed and preceded by a less-than symbol (<). Any lines that exist in the current, but not in the old will be listed and preceded by a greater-than symbol (>). Lines that are the same in both files are not listed; for example, files that are identical will produce no output from the diff command. If the location or description changes, lines preceded by both < and > will be output.

    If the system is configured with more than one logical partition with Linux installed, repeat "14" on page 52 and "15" on page 52 for all logical partitions that have Linux installed.

16. Was the location of one and only one device recorded in "15" on page 52?

    No

    If you previously answered Yes to step "16" on page 52, return the system to its original configuration. **This ends the procedure**.

    Go to MAP 0410: Repair checkout.

    If you did not previously answer Yes to step "16" on page 52, go to step "17" on page 53.

    Yes

    Complete the following steps one at a time. Power off the system before each step. After each step, power on the system and go to step "13" on page 52.

    a. Check all connections from the system to the device.

    b. Replace the device (for example, tape or DASD).

    c. If applicable, replace the device backplane.

    d. Replace the device cable.

e. Replace the adapter.

   - If the adapter resides in an I/O drawer, replace the I/O backplane.
   - If the device adapter resides in the CEC, replace the I/O riser card, or the CEC backplane in which the adapter is plugged.

f. Call service support. Do not go to step “13” on page 52.

17. Does the system appear to stop or hang before reaching the login prompt or did you record any problems with resources in step “15” on page 52?

    Note: If the system console or VTERM window is always blank, choose NO. If you are sure the console or VTERM is operational and connected correctly, answer the question for this step.

    No
    Go to step “18” on page 53.

    Yes
    There may be a problem with an I/O device. Go to PFW1542: I/O problem isolation procedure. When instructed to boot the system, boot a full system partition.

18. Boot the eServer standalone diagnostics, refer to Running the online and stand-alone diagnostics.

    Run diagnostics in problem determination mode on all resources. Be sure to boot a full system partition. Ensure that diagnostics were run on all known resources. You may need to select each resource individually and run diagnostics on each resource one at a time.

    Did standalone diagnostics find a problem?

    No
    Go to step “22” on page 53.

    Yes
    Go to the Reference codes and perform the actions for each reference code you have recorded. For each reference code not already processed in step “16” on page 52, repeat this action until the problem has been corrected. Perform the indicated actions, one at a time. If all recorded reference codes have been processed and the problem has not been corrected, go to step “22” on page 53.

19. Does the system have Linux installed on one or more partitions?

    No
    Return to the “Beginning problem analysis” on page 1.

    Yes
    Go to step “3” on page 50.

20. Were any location codes recorded in steps “3” on page 50, “6” on page 50, “8” on page 51, “9” on page 51, “10” on page 51, or “11” on page 51?

    No
    Go to step “13” on page 52.

    Yes
    Replace, one at a time, all parts whose location code was recorded in steps “3” on page 50, “6” on page 50, “8” on page 51, “9” on page 51, “10” on page 51, or “11” on page 51 that have not been replaced. Power off the system before replacing a part. After replacing the part, power on the system to check if the problem has been corrected. Go to step “21” on page 53 when the problem has been corrected, or all parts in the location codes list have been replaced.

21. Was the problem corrected in step “20” on page 53?

    No
    Go to step “13” on page 52.

    Yes
    Return the system to its original configuration. **This ends the procedure**.

    Go to MAP 0410: Repair checkout.

22. Were any other symptoms recorded in step “3” on page 50?

No

Call support.

Yes

Go to the "Beginning problem analysis" on page 1 with each symptom you have recorded. Perform the indicated actions for all recorded symptoms, one at a time, until the problem has been corrected. If all recorded symptoms have been processed and the problem has not been corrected, call your next level of support.

Detecting problems

Provides information on using various tools and techniques to detect and identify problems.

IBM i problem determination procedures

There are several tools you can use to determine a problem with an IBM i system or partition.

For information about these tools, see the following topics:

- Searching the service action log
- Using the product activity log
-

Using the problem log

Use this procedure to find and analyze a problem log entry that relates to the problem reported.

About this task

Note: For on-line problem analysis (WRKPRB), ensure that you are logged on with QSRV authority. During problem isolation, this will allow access to test procedures that are not available under any other log-on.

Procedure

1. On the command line, enter the Work with Problems command:

   ```
   WRKPRB
   ```

   Note: Use F4 to change the WRKPRB parameters to select and sort on specific problem log entries that match the problem. Also, F11 displays the dates and times the problems were logged by the system.

   Was an entry that relates to the problem found?

   Note: If the WRKPRB function was not available answer NO.

   Yes: Continue with the next step.
   No: Go to Problems with noncritical resources. **This ends the procedure**.
2. Select the problem entry by moving the cursor to the problem entry option field and entering option 8 to work with the problem.

   Is Analyze Problem (option 1) available on the Work with Problem display?

   No: Perform the following:

   a. Return to the initial problem log display (F12).
   b. Select the problem entry by moving the cursor to the problem entry option field and selecting the option to display details.
   c. Select the function key to display possible causes.

      Note: If this function key is not available, use the customer reported symptom string for customer perceived information about this problem. Then, go to Using the product activity log.
   d. Use the list of possible causes as the FRU list and go to step "5" on page 55.

Yes: Run Analyze Problem (option 1) from the Work with Problem display.

Notes:

a. For SRCs starting with 6112 or 9337, use the SRC and go to the Reference codes topic.
b. If the message on the display directs you to use SST (System Service Tools), go to COMIP01.

Was the problem corrected by the analysis procedure?

No: Continue with the next step.
Yes: **This ends the procedure**.

3. Did problem analysis send you to another entry point in the service information?

   No: Continue with the next step.
   Yes: Go to the entry point indicated by problem analysis. **This ends the procedure**.

4. Was the problem isolated to a list of failing items?

   Yes: Continue with the next step.
   No: Go to Problems with noncritical resources. **This ends the procedure**.

5. Exchange the failing items one at a time until the problem is repaired.

   Notes:

   a. For Symbolic FRUs, see Symbolic FRUs.
   b. When exchanging FRUs, go to the remove and replace procedures for your specific system.

   Has the problem been resolved?

   No: Contact your next level of support. **This ends the procedure**.
   Yes: **This ends the procedure**.

Problem determination procedure for AIX or Linux servers or partitions

This procedure helps to produce or retrieve a service request number (SRN) if the customer or a previous procedure did not provide one.

If your server is running AIX or Linux, use one of the following procedures to test the server or partition resources to help you determine where a problem might exist.

If you are servicing a server running the AIX operating system, go to MAP 0020: Problem determination procedure.

If you are servicing a server running the Linux operating system, go to the Linux problem isolation procedure.

System unit problem determination

Use this procedure to obtain a reference code if the customer did not provide you with one, or you are unable to load server diagnostics.

If you are able to load the diagnostics, go to Problem determination procedure for AIX or Linux servers or partitions.

The service processor may have recorded one or more symptoms in its error log. Examine this error log before proceeding (For more information, see Managing your server using the Advanced System Management Interface). The server may have been set up by using the management console. Check the Service Action Event (SAE) log in the Service Focal Point. The SAE log may have recorded one or more symptoms in the Service Focal Point. To avoid unnecessary replacement of the same FRU for the same problem, it is necessary to check the SAE log for evidence of prior service activity on the same subsystem.

The service processor may have been set by the user to monitor system operations and to attempt recoveries. You can disable these actions while you diagnose and service the system. If the system maintenance policies were saved by using the save/restore hardware maintenance policies, all the settings of the service processor (except language) were saved and you can use the same service aid to restore the settings at the conclusion of your service action.

If you disable the service processor settings, note their current settings so that you can restore when you are done.

If the system is set to power on using one of the parameters in the following table, disconnect the modem to prevent incoming signals that could cause the system to power on.

Following are the service processor settings. For more information about the service processor settings, see Managing your server using the Advanced System Management Interface.

Table 10. Service processor settings

Setting	Description
Monitoring (also called surveillance)	From the ASMI menu, expand the **System Configuration**, then click on **Monitoring**. Disable both types of surveillance.
Auto power restart (also called unattended start mode)	From the ASMI menu, expand **Power/Restart Control**, then click on **Auto Power Restart**, and set it to disabled.
Wake on LAN	From the ASMI menu, expand **Wake on LAN**, and set it to disabled
Call out	From the ASMI menu, expand the **Service Aids**, then click on **Call-Home/Call-In Setup**. Set the call-home system port and the call-in system port to disabled.

Step 1020-1

Be prepared to record code numbers to help analyze a problem.

Analyze a failure to load the diagnostic programs

Follow these steps to analyze a failure to load the diagnostic programs.

Note: Be prepared to answer questions regarding the control panel and to perform certain actions based on displayed POST indicators. Observer these conditions.

1. Run diagnostics on any partition. Find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, continue to the next step.
2. Run diagnostics on the failing partition. Find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, continue to the next step.
3. Power off the system.
4. Load the standalone diagnostics in service mode to test the full system partition. For more information, see Running the online and stand-alone diagnostics.
5. Wait until the diagnostics are loaded or the system appears to stop. If you receive an error code or if the system stops before diagnostics are loaded, find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, continue to the next step.
6. Run the standalone diagnostics on the entire system. Find your symptom in the following table, then follow the instructions given in the Action column. If no fault is identified, call service support for assistance.

Symptom	Action
One or more logical partitions does not boot.	a. Check service processor error log. If an error is indicated, go to "Beginning problem analysis" on page 1. b. Check the Serviceable action event log, go to "Beginning problem analysis" on page 1. c. Go to Problems with loading and starting the operating system.
The rack identify LED does not operate properly.	Go to "Beginning problem analysis" on page 1.
The system stopped and a system reference code is displayed on the operator panel.	Go to "Beginning problem analysis" on page 1.
The system stops with a prompt to enter a password.	Enter the password. You cannot continue until a correct password has been entered. When you have entered a valid password, go to the beginning of this table and wait for one of the other conditions to occur.
The diagnostic operating instructions are displayed.	Go to MAP 0020: AIX or Linux problem determination procedure.
The power good LED does not come on or does not stay on, or you have a power problem.	Go to Power problems.
The system login prompt is displayed.	You may not have pressed the correct key or you may not have pressed the key soon enough when you were to trying to indicate a service mode IPL of the diagnostic programs. If this is the case, start again at the beginning of this step. **Note:** Perform the system shutdown procedure before turning off the system. If you are sure you pressed the correct key in a timely manner, go to Step 1020-2.
The system does not respond when the password is entered.	Go to Step 1020-2.

Symptom	Action
The system stopped. A POST indicator is displayed on the system console and an eight-digit error code is not displayed.	If the POST indicator represents: a. Memory, go to PFW 1548: Memory and processor subsystem problem isolation procedure. b. Keyboard i) Replace the keyboard cable. ii) Replace the keyboard. iii) Replace the service processor. Location is model dependent. iv) Go to PFW1542: I/O problem isolation procedure. c. Network, go to PFW1542: I/O problem isolation procedure. d. SCSI, go to PFW1542: I/O problem isolation procedure. e. Speaker i) Replace the operator panel. Location is model dependent. ii) Replace the service processor. Location is model dependent. iii) Go to PFW1542: I/O problem isolation procedure.
The System Management Services menu is displayed.	Go to PFW1542: I/O problem isolation procedure.
All other symptoms.	If you were directed here from the Entry MAP, go to PFW1542: I/O problem isolation procedure. Otherwise, find the symptom in the "Beginning problem analysis" on page 1.

Step 1020-2

Use this procedure to analyze a keyboard problem.

Find the type of keyboard you are using in the following table; then follow the instructions given in the Action column.

Keyboard Type	Action
Type 101 keyboard (U.S.). Identified by the size of the Enter key. The Enter key is in only one horizontal row of keys.	Record error code M0KB D001; then go to Step 1020-3.
Type 102 keyboard (W.T.). Identified by the size of the Enter key. The Enter key extends into two horizontal rows.	Record error code M0KB D002; then go to Step 1020-3.
Type 106 keyboard. (Identified by the Japanese characters.)	Record error code M0KB D003; then go to Step 1020-3.
ASCII terminal keyboard	Go to the documentation for this type of ASCII terminal and continue with problem determination.

Step 1020-3

Perform the following steps:

1. Find the 8-digit error code in Reference codes.

   Note: If you do not locate the 8-digit code, look for it in one of the following places:

- Any supplemental service manuals for attached devices
- The diagnostic problem report screen for additional information
- The Service Hints service aid
- The CEREADME file

2. Perform the action listed.

Management console machine code problems

The support organization uses the *pesh* command to look at the management console internal machine code to determine how to fix a machine code problem. Only a service representative or support representative can access this feature.

Launching an xterm shell

About this task

You may need to launch an xterm shell to perform directed support from the support center. This may be required if the support center needs to analyze a system dump in order to better understand machine code operations at the time of a failure. To launch an xterm shell, perform the following:

Procedure

1. Open a terminal by right-clicking the background and selecting **Terminals** > **rshterm**.
2. Type the *pesh* command followed by the serial number of the management console and press Enter.
3. You will be prompted for a password, which you must obtain from your next level of support.

Results

Additional information: "Viewing the management console logs" on page 59.

Viewing the management console logs

The console logs display error and information messages that the console has logged while running commands.

About this task

The service representative can use this information to learn more about what caused an error and how to resolve it. The management console classifies log entries as either an informational message or an error message. Log entries are identified with an *I* or *E*, respectively. The management console lists these log entries chronologically, with the most recent shown at the top of the list.

Use the management console Log to view a record of management console system events. System events are activities that indicate when processes begin and end. These events also indicate whether the attempted action was successful.

To view the HMC log, perform the following:

1. Launch an xterm shell (see "Launching an xterm shell" on page 59).
2. When you have entered the password, use the *showLog* command to launch the HMC log window.

The log includes the following information:

- The event's unique ID code
- The date the event occurred
- The time the event occurred
- The log's type
- The name of the attempted action

- The log's reference code
- The status of the log

View a particular event

About this task

To view a particular event, perform the following steps:

Procedure

1. Select an event by clicking once on it.
2. Press Enter to get to a summary of the log you selected. From here, you must select a Block ID to display. The blocks are listed next to the buttons and include the following options:
   - Standard Data Block
   - Secondary Data Block
   - Microcode Reason / ID Error Information
3. Select the data block you want to view.
4. Press Enter. The extended information shown for the data blocks you selected includes the following items:
   - Program name
   - Current process ID
   - Parent process ID
   - Current thread priority
   - Current thread ID
   - Screen group
   - Subscreen group
   - Current foreground screen process group
   - Current background screen process group

Problem determination procedures

Problem determination procedures are provided by power-on self-tests (POSTs), service request numbers, and maintenance analysis procedures (MAPs). Some of these procedures use the service aids that are described in the user or maintenance information for your system SCSI attachment.

Disk drive module power-on self-tests

The disk drive module Power-on Self-Tests (POSTs) start each time that the module is switched on, or when a Send Diagnostic command is received. They check whether the disk drive module is working correctly. The POSTs also help verify a repair after a Field Replaceable Unit (FRU) has been exchanged.

The tests are POST-1 and POST-2.

POST-1 runs immediately after the power-on reset line goes inactive, and before the disk drive module motor starts. POST-1 includes the following tests:

- Microprocessor
- ROM
- Checking circuits

If POST-1 completes successfully, POST-2 is enabled.

If POST-1 fails, the disk drive module is not configured into the system.

POST-2 runs after the disk drive module motor has started. POST-2 includes the following tests:

- Motor control
- Servo control
- Read and write on the diagnostic cylinder (repeated for all heads)
- Error checking and correction (ECC).

If POST-2 completes successfully, the disk drive module is ready for use with the system.

If POST-2 fails, the disk drive module is not configured into the system.

SCSI card power-on self-tests

The SCSI card Power-On Self Tests (POSTs) start each time power is switched on, or when a Reset command is sent from the using system SCSI attachment. They check only the internal components of the SCSI card; they do not check any interfaces to other FRUs.

If the POSTs complete successfully, control passes to the functional microcode of the SCSI card. This microcode checks all the internal interfaces of the I/O enclosure, and reports failures to the host system.

If the POSTs fail, one of the following events occur:

- The SCSI card check LED and the enclosure check LED come on.
- If the SCSI was configured for high availability using a dual initiator card the error will be reported. However, the functional operation of the enclosure is not affected. For example, the customer still has access to all the disk drive modules.

The failure is reported when:

- the failure occurs at system bring-up time, the host system might detect that the enclosure is missing, and reports an error.
- the failure occurs at any time other than system bring-up time, the hourly health check reports the failure.

Analyzing problems

Use these instructions and procedures to help you determine the cause of the problem.

Intermittent problems

An intermittent problem is a problem that occurs for a short time, and then goes away.

About this task

The problem may not occur again until some time in the future, if at all. Intermittent problems cannot be made to appear again easily.

Some examples of intermittent problems are:

- A reference code appears on the control panel (the system attention light is on) but disappears when you power off, then power on the system. An entry does not appear in the Product Activity Log.
- An entry appears in the problem log when you use the Work with Problems (`WRKPRB`) command. For example, an expansion unit becomes powered off, but starts working again when you power it on.
- The workstation adapter is in a hang condition but starts working normally when it gets reset.

Note: You can get equipment for the following conditions from your branch office or installation planning representative:

- If you suspect that the air at the system site is too hot or too cold, you need a thermometer to check the temperature.
- If you suspect the moisture content of the air at the system site is too low or too high, use a wet/dry bulb to check the humidity. See "General intermittent problem checklist" on page 63 for more information.

- If you need to check AC receptacles for correct wiring, you need an ECOS tester, Model 1023-100, or equivalent tester. The tester lets you quickly check the receptacles. If you cannot find a tester, use an analog multimeter instead. Do not use a digital multimeter.

Follow the steps below to correct an intermittent problem:

Procedure

1. Read the information in "About intermittent problems" on page 62 before you attempt to correct an intermittent problem.

   Then, continue with the next step of this procedure.
2. Perform *all* steps in the "General intermittent problem checklist" on page 63.

   Then, continue with the next step of this procedure.
3. Did you correct the intermittent problem?

   Yes: This ends the procedure.

   No: Go to "Analyzing intermittent problems" on page 65. **This ends the procedure.**

About intermittent problems

An intermittent problem can show many different symptoms, so it might be difficult for you to determine the real cause without completely analyzing the failure.

To help with this analysis, you should determine as many symptoms as possible.

- The complete reference code is necessary to determine the exact failing area and the probable cause.
- Product activity log (PAL) information can provide time and device relationships.
- Information about environmental conditions when the failure occurred can be helpful (for example, an electrical storm occurring when the failure happened).

Note: If you suspect that an intermittent problem is occurring, increase the log sizes to the largest sizes possible. Select the PAL option on the Start a Service Tool display (see Using the product activity log for details).

Types of intermittent problems

Following are the major types of intermittent problems:

- Code (PTFs):
  - Licensed Internal Code
  - IBM i
  - Licensed program products
  - Other application software
- Configuration:
  - Non-supported hardware that is used on the system
  - Non-supported system configurations
  - Non-supported communication networks
  - Model and feature upgrades that are not performed correctly
  - Incorrectly configured or incorrectly cabled devices
- Environment:
  - Power line disturbance (for example, reduced voltage, a pulse, a surge, or total loss of voltage on the incoming AC voltage line)
  - Power line transient (for example, lightning strike)
  - Electrical noise (constant or intermittent)

  - Defective grounding or a ground potential difference
  - Mechanical vibration
- Intermittent hardware failure

General intermittent problem checklist

Use the following procedure to correct intermittent problems.

About this task

Performing these steps removes the known causes of most intermittent problems.

Procedure

1. Discuss the problem with the customer.

   Look for the following symptoms:

   - A reference code that goes away when you power off and then power on the system.
   - Repeated failure patterns that you cannot explain. For example, the problem occurs at the same time of day or on the same day of the week.
   - Failures that started after system relocation.
   - Failures that occurred during the time specific jobs or software were running.
   - Failures that started after recent service or customer actions, system upgrade, addition of I/O devices, new software, or program temporary fix (PTF) installation.
   - Failures occurring only during high system usage.
   - Failures occur when people are close to the system or machines are attached to the system.
2. Recommend that the customer install the latest cumulative PTF package, since code PTFs have corrected many problems that seem to be hardware failures.

   The customer can order the latest cumulative PTF package electronically through Electronic Customer Support or by calling the Software Support Center.
3. If you have not already done so, use the maintenance package to see the indicated actions for the symptom described by the customer.

   Attempt to perform the on-line problem analysis procedure first. If this is not possible, such as when the system is down, go to the “Beginning problem analysis” on page 1.

   Use additional diagnostic tools, if necessary, and attempt to recreate the problem.

   Note: Ensure that the service information you are using is at the same level as the operating system.
4. Check the site for the following environmental conditions:

   a) Any electrical noise that matches the start of the intermittent problems. Ask the customer such questions as:

   - Have any external changes or additions, such as building wiring, air conditioning, or elevators been made to the site?
   - Has any arc welding occurred in the area?
   - Has any heavy industrial equipment, such as cranes, been operating in the area?
   - Have there been any thunderstorms in the area?
   - Have the building lights become dim?
   - Has any equipment been relocated, especially computer equipment?

   If there was any electrical noise, find its source and prevent the noise from getting into the system.

   b) Site temperature and humidity conditions that are within the system specifications.

   See temperature and humidity design criteria in the Planning for the system topic relevant for your system.

   c) Poor air quality in the computer room:

      - Look for dust on top of objects. Dust particles in the air cause poor electrical connections and may cause disk unit failures.
      - Smell for unusual odors in the air. Some gases can corrode electrical connections.

   d) Any large vibration (caused by thunder, an earthquake, an explosion, or road construction) that occurred in the area at the time of the failure.

      Note: A failure that is caused by vibration is more probable if the server is on a raised floor.

5. Ensure that all ground connections are tight.

   These items reduce the effects of electrical noise. Check the ground connections by measuring the resistance between a conductive place on the frame to building ground or to earth ground. The resistance must be 1.0 ohm or less.

6. Ensure proper cable retention is used, as provided.

   If no retention is provided, the cable should be strapped to the frame to release tension on cable connections.

   Ensure that you pull the cable ties tight enough to fasten the cable to the frame bar tightly. A loose cable can be accidentally pulled with enough force to unseat the logic card in the frame to which the cable is attached. If the system is powered on, the logic card could be destroyed.

7. Ensure that all workstation and communications cables meet hardware specifications:

   - All connections are tight.
   - Any twinaxial cables that are not attached to devices must be removed.
   - The lengths and numbers of connections in the cables must be correct.
   - Ensure that lightning protection is installed on any twinaxial cables that enter or leave the building.

8. Perform the following:

   a) Review recent repair actions.

      Contact your next level of support for assistance.

   b) Review entries in the problem log (WRKPRB).

      Look for problems that were reported to the user.

   c) Review entries in the PAL, SAL, and service processor log. Look for a pattern:

      - SRCs on multiple adapters occurring at the same time
      - SRCs that have a common time-of-day or day-of-week pattern
      - Log is wrapping (hundreds of recent entries and no older entries)

      Check the PAL sizes and increase them if they are smaller than recommended.

   d) Review entries in the history log (`Display Log (DSPLOG)`).

      Look for a change that matches the start of the intermittent problems.

   e) Ensure that the latest engineering changes are installed on the system and on all system I/O devices.

9. Ensure that the hardware configuration is correct and that the model configuration rules have been followed.

   Use the **Display hardware configuration** service function (under SST or DST) to check for any missing or failed hardware.

10. Was a system upgrade, feature, or any other field bill of material or feature field bill of material installed just before the intermittent problems started occurring?

    No: Continue with the next step.

    Yes: Review the installation instructions to ensure that each step was performed correctly. Then, continue with the next step of this procedure.

11. Is the problem associated with a removable media storage device?

No: Continue with the next step.

Yes: Ensure that the customer is using the correct removable media storage device cleaning procedures and good storage media. Then, continue with the next step of this procedure.

12. Perform the following to help prevent intermittent thermal checks:

    - Ensure that the AMDs are working.
    - Exchange all air filters as recommended.

13. If necessary, review the intermittent problems with your next level of support and installation planning representative.

    Ensure that all installation planning checks were made on the system. Because external conditions are constantly changing, the site may need to be checked again. **This ends the procedure.**

Analyzing intermittent problems

This procedure enables you to begin analyzing an intermittent problem.

About this task

Use this procedure only after you have first reviewed the information in "About intermittent problems" on page 62 and gone through the "General intermittent problem checklist" on page 63.

Procedure

1. Is a reference code associated with the intermittent problem?

   No: Continue with the next step.

   Yes: Go to Reference codes. If the actions in the reference code tables do not correct the intermittent problem, return here and continue with the next step.

2. Is a symptom associated with the intermittent problem?

   No: Continue with the next step.

   Yes: Go to "Intermittent symptoms" on page 65. If the information there does not help to correct the intermittent problem, return here and continue with the next step.

3. Go to "Failing area intermittent isolation procedures" on page 66.

   If the information there does not help to correct the intermittent problem, return here and then, continue with the next step.

4. Send the data you have collected to your next level of support so that an Authorized Program Analysis Report (APAR) can be written.

   This ends the procedure.

Intermittent symptoms

Use the table below to find the symptom and description of the intermittent problem. Then perform the corresponding intermittent isolation procedures.

Although an isolation procedure may correct the intermittent problem, use your best judgment to determine if you should perform the remainder of the procedure shown for the symptom.

Note: If the symptom for the intermittent problem you have is not listed, go to "Failing area intermittent isolation procedures" on page 66.

Table 11. Intermittent symptoms

Symptom	Description	Isolation procedure
System powered off.	The system was operating correctly, then the system powered off. A 1xxx SRC may occur when this happens, and this SRC info should be logged in the service processor log.	INTIP09

Table 11. Intermittent symptoms (continued)

Symptom	Description	Isolation procedure
System stops.	The system is powered on but is not operating correctly. No SRC is displayed. The system attention light is off and the processor activity lights may be on or off. Noise on the power-on reset line can cause the processor to stop.	INTIP18
System or subsystem runs slow.	The system or the subsystem is not processing at its normal speed.	INTIP20

Failing area intermittent isolation procedures

This procedure helps you determine how to resolve intermittent problems when you do not have a system reference code (SRC) or cannot determine the symptom.

About this task

Use this table only if you do not have a system reference code (SRC), or cannot find your symptom in "Intermittent symptoms" on page 65.

Procedure

1. Perform all of the steps in "General intermittent problem checklist" on page 63 for all failing areas.

   Then, continue with the next step.
2. Refer to the table below, and perform the following:
   a) Find the specific area of failure under **Failing area**.
   b) Look down the column of the area of failure until you find an X.
   c) Look across to the **Isolation procedure** column and perform the procedure indicated.
   d) If the isolation procedure does not correct the intermittent problem, continue down the column of the area of failure until you have performed all of the procedures shown for the failing area.
3. Although an isolation procedure may correct the intermittent problem, use your best judgment to determine if you should perform the remainder of the procedures shown for the failing area.

Results

Table 12. Failing area intermittent isolation procedures

Failing area						Isolation procedure to perform:
Power	**Work station I/O processor**	**Disk unit adapter**	**Comm-unication**	**Processor bus**	**Tape optical**	**Perform all steps in:**
X	X	X	X	X	X	"General intermittent problem checklist" on page 63
X	X			X		INTIP05
	X	X	X	X	X	INTIP07
X						INTIP09
X						INTIP14

Table 12. Failing area intermittent isolation procedures (continued)

Failing area						**Isolation procedure to perform:**
Power	**Work station I/O processor**	**Disk unit adapter**	**Comm-unication**	**Processor bus**	**Tape optical**	**Perform all steps in:**
		X				INTIP16
X	X	X	X	X	X	INTIP18
	X	X	X	X	X	INTIP20

IPL problems

Use these scenarios to help you diagnose your IPL problem.

Cannot perform IPL from the control panel (no SRC)

Use this procedure when you cannot perform an IBM i IPL from the control panel (no SRC).

About this task

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

Procedure

1. Perform the following:
   a) Verify that the power cable is plugged into the power outlet.
   b) Verify that power is available at the customer's power outlet.
2. Start an IPL by doing the following:
   a) Select Manual mode and IPL type A or B on the control panel. See Control panel functions for details.
   b) Power on the system. See Starting a system.

   Does the IPL complete successfully?

   No: Continue with the next step.
   Yes: **This ends the procedure.**
3. Have all the units in the system become powered on that you expected to become powered on?

   Yes: Continue with the next step.
   No: Go to Power problems and find the symptom that matches the problem. **This ends the procedure.**
4. Is an SRC displayed on the control panel?
   - **Yes**: Go to Power problems and use the displayed SRC to correct the problem. **This ends the procedure.**
   - **No**: For all models, exchange the following FRUs, one at a time. Refer to the remove and replace procedures for your specific system for additional information.

   a) SPCN card unit. See symbolic FRU TWRCARD.
   b) Power Supply. See symbolic FRU PWRSPLY. **This ends the procedure.**

Cannot perform IPL at a specified time (no SRC)

Use this procedure when you cannot perform an IBM i IPL at a specified time (no SRC). To correct the IPL problem, perform this procedure until you determine the problem and can perform an IPL at a specified time.

About this task

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

Procedure

1. Verify the following:
   a) The power cable is plugged into the power outlet.
   b) That power is available at the customer's power outlet.
2. Power on the system in normal mode. See Starting a system.

   Does the IPL complete successfully?

   Yes: Continue with the next step.
   No: Go to the “Beginning problem analysis” on page 1 procedure. **This ends the procedure.**
3. Have all the units in the system become powered on that you expected to become powered on?

   Yes: Continue with the next step.
   No: Go to “Beginning problem analysis” on page 1 and find the symptom that matches the problem. **This ends the procedure.**
4. Verify the requested system IPL date and time by doing the following:
   a) On the command line, enter the Display System Value command:

   ```
   DSPSYSVAL QIPLDATTIM
   ```

   Observe the system value parameters.

   Note: The system value parameters are the date and time the system operator requested a timed IPL.

   ```
   +------------------------------------------------------------------------------+
   |Display System Value                                                          |
   |System:  S0000000                                                             |
   |System value . . . . . . . . . :  QIPLDATTIM                                  |
   |                                                                              |
   |Description  . . . . . . . . . :  Date and time to automatically IPL          |
   |                                                                              |
   |                                                                              |
   |IPL date     . . . . . . . . . :  MM/DD/YY                                    |
   |IPL time     . . . . . . . . . :  HH:MM:SS                                    |
   +------------------------------------------------------------------------------+
   ```

   Figure 1. Display for QIPLDATTIM

   b) Verify the system date. On the command line, enter the Display System Value command:

   ```
   DSPSYSVAL QDATE
   ```

   Check the system values for the date.

```
+------------------------------------------------------------------------------+
|Display System Value                                                          |
|System:  S0000000                                                             |
|System value . . . . . . . . . . :  QDATE                                     |
|                                                                              |
|Description  . . . . . . . . . . :  System date                               |
|                                                                              |
|Date          . . . . . . . . . . :  MM/DD/YY                                 |
+------------------------------------------------------------------------------+
```

Figure 2. Display for QDATE

Does the operating system have the correct date?

- **Yes**: Continue with this step.
- **No**: Set the correct date by doing the following:

  i) On the command line, enter the Change System Value command (`CHGSYSVAL QDATE VALUE('mmddyy')`).

  ii) Set the date by entering

      mm=month
      dd=day
      yy=year

  iii) Press **Enter**.

c) Verify the system time. On the command line, enter the Display System Value command: `DSPSYSVAL QTIME`

Check the system values for the time.

```
+------------------------------------------------------------------------------+
|Display System Value                                                          |
|System:  S0000000                                                             |
|System value . . . . . . . . . . :  QTIME                                     |
|                                                                              |
|Description  . . . . . . . . . . :  Time of day                               |
|                                                                              |
|Time          . . . . . . . . . . :  HH:MM:SS                                 |
+------------------------------------------------------------------------------+
```

Figure 3. Display for QTIME

Does the operating system have the correct time?

- **Yes**: Continue with this step.
- **No**: Set the correct time by doing the following:

  i) On the command line, enter the Change System Value command (`CHGSYSVAL QTIME VALUE('hhmmss')`).

  ii) Set the time by entering

      hh=24 hour time clock
      mm=minutes
      ss=seconds

  iii) Press **Enter** and then, continue with the next step.

5. Verify that the system can perform an IPL at a specified time by doing the following:

   a) Set the IPL time to 5 minutes past the present time by entering the Change System Value command (`CHGSYSVAL SYSVAL(QIPLDATTIM) VALUE('mmddyy hhmmss')`) on the command line.

      mm = month to power on
      dd = day to power on
      yy = year to power on
      hh = hour to power on

mm = minute to power on
ss = second to power on

b) Power off the system by entering the Power Down System Immediate command (`PWRDWNSYS *IMMED`) on the command line.

c) Wait 5 minutes.

Does the IPL start at the time you specified?

No: Continue with the next step.
Yes: **This ends the procedure.**

6. Power on the system in normal mode. See Starting a system.

   Does the IPL complete successfully?

   Yes: Continue with the next step.
   No: Go to “Beginning problem analysis” on page 1. **This ends the procedure.**

7. Find an entry in the Service Action Log that matches the time, SRC, and/or resource that compares to the reported problem.

   a) On the command line, enter the Start System Service Tools command:

   ```
   STRSST
   ```

   If you cannot get to SST, select DST. See Dedicated service tools (DST) for details.

   Note: Do not IPL the system or partition to get to DST.

   b) On the Start Service Tools Sign On display, type in a user ID with service authority and password.

   c) Select **Start a Service Tool** > **Hardware Service Manager** > **Work with service action log**.

   d) On the Select Timeframe display, change the From: Date and Time to a date and time prior to when the customer reported having the problem.

   e) Find an entry that matches one or more conditions of the problem:

   - SRC
   - Resource
   - Time
   - FRU list (choose **Display the failing item information** to display the FRU list).

   Notes:

   a. All entries in the service action log represent problems that require a service action. It may be necessary to handle any problem in the log even if it does not match the original problem symptom.

   b. The information displayed in the date and time fields are the time and date for the first occurrence of the specific system reference code (SRC) for the resource displayed during the time range selected.

   Did you find an entry in the Service Action Log?

   No: Continue with the next step.
   Yes: Go to step “9” on page 71.

8. Exchange the following parts one at a time.

   See the remove and replace procedures for your specific system. After exchanging each part, return to step “5” on page 69 to verify that the system can perform an IPL at a specified time.

   Note: If you exchange the control panel or the system backplane, you must set the correct date and time by performing step “4” on page 68.

   ⚠ **Attention:** Before exchanging any part, power off the system. See Stopping a system.

   - System unit backplane (see symbolic FRU SYSBKPL)

- System control panel
- System control panel cable

Did the IPL complete successfully after you exchanged all of the parts listed above?

No: Contact your next level of support. **This ends the procedure.**
Yes: Continue with the next step.

9. Was the entry isolated (is there a Y in the Isolated column)?

   - **No**: Go to Reference codes and use the SRC indicated in the log. **This ends the procedure.**
   - **Yes**: Display the failing item information for the Service Action Log entry. Items at the top of the failing item list are more likely to fix the problem than items at the bottom of the list.

     Exchange the failing items one at a time until the problem is repaired. After exchanging each one of the items, verify that the item exchanged repaired the problem.

     Notes:

     a. For symbolic FRUs see Symbolic FRUs.
     b. When exchanging FRUs, refer to the remove and replace procedures for your specific system.
     c. After exchanging an item, go to Verifying the repair.

     After the problem has been resolved, close the log entry by selecting **Close a NEW entry** on the Service Actions Log Report display. **This ends the procedure.**

Cannot automatically perform an IPL after a power failure

Use this procedure when you cannot automatically perform an IBM i IPL after a power failure.

Procedure

1. Normal or Auto mode on the control panel must be selected when power is returned to the system.

   Is Normal or Auto mode on the control panel selected?

   Yes: Continue with the next step.
   No: Select **Normal** or **Auto** mode on the control panel. **This ends the procedure.**

2. Use the Display System Value command (`DSPSYSVAL`) to verify that the system value under QPWRRSTIPL on the Display System Value display is equal to 1.

   Is QPWRRSTIPL equal to 1?

   Yes: Contact your next level of support.
   No: Use the Change System Value command (`CHGSYSVAL`) to set QPWRRSTIPL equal to 1. **This ends the procedure.**

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with ICT Accessibility 508 Standards and 255 Guidelines (https://www.access-board.gov/ict/) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power servers.

The IBM Power servers online product documentation in IBM Documentation is enabled for accessibility. For more information about IBM's commitment to accessibility, see the IBM accessibility website at IBM Accessibility (https://www.ibm.com/able/).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Electronic emission notices

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Class A Notices

The following Class A statements apply to the IBM servers that contain the Power10 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

The following Class A statements apply to the servers.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) ". Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

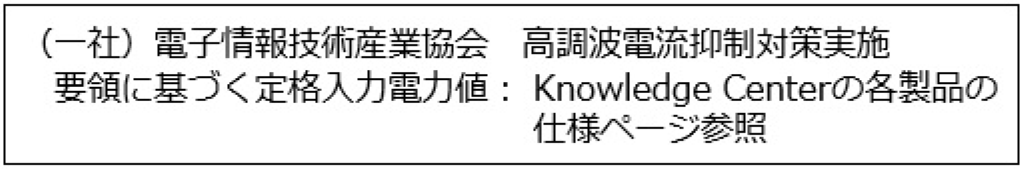

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

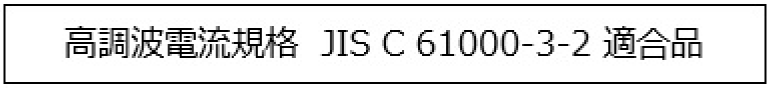

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

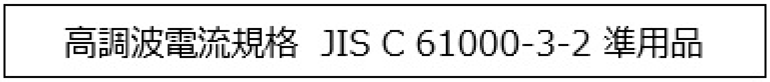

高調波電流規格 JIS C 61000-3-2 準用品

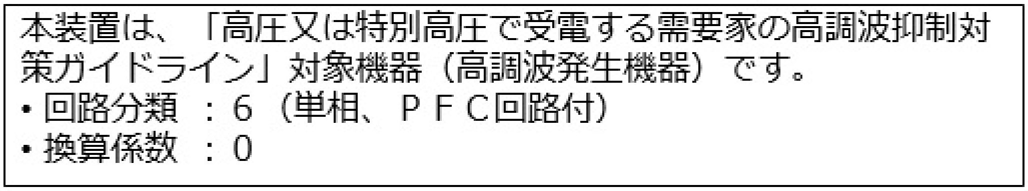

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：6（単相、ＰＦＣ回路付）
- 換算係数 ：0

This statement applies to products greater than 20 A per phase, three-phase.

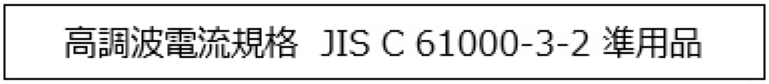

高調波電流規格 JIS C 61000-3-2 準用品

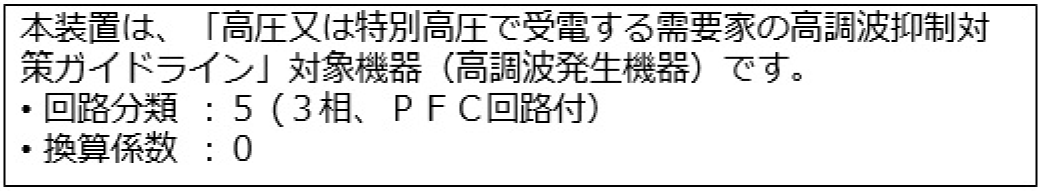

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：5（3相、ＰＦＣ回路付）
- 換算係数 ：0

Japan Voluntary Control Council for Interference (VCCI) Notice

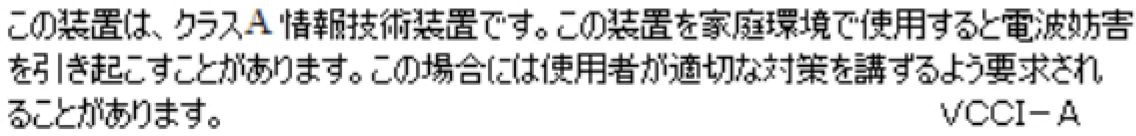

この装置は、クラスA 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。　VCCI－A

Korea Notice

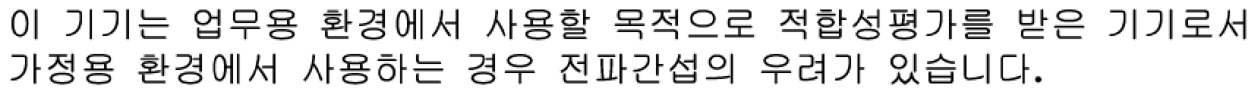

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서
가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

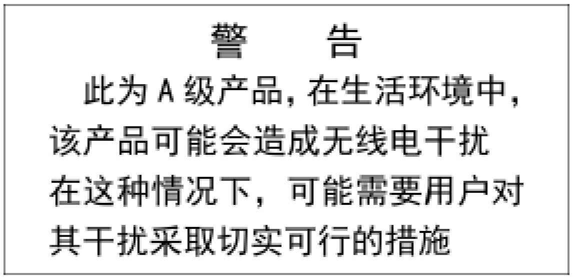

警　告

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰
在这种情况下，可能需要用户对
其干扰采取切实可行的措施

Russia Notice

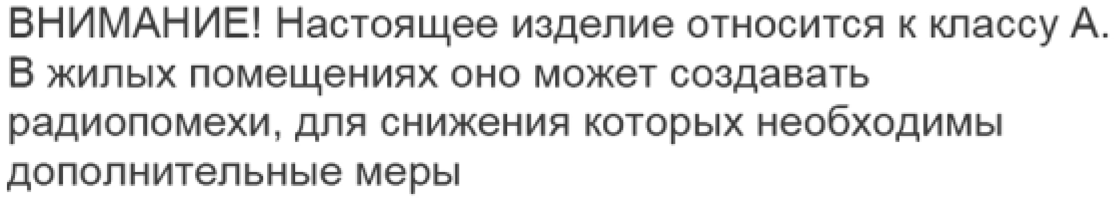

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

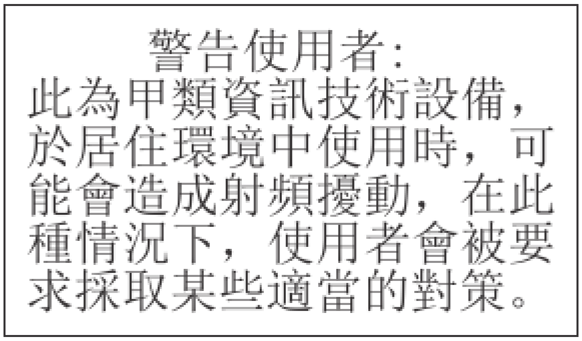

警告使用者：
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

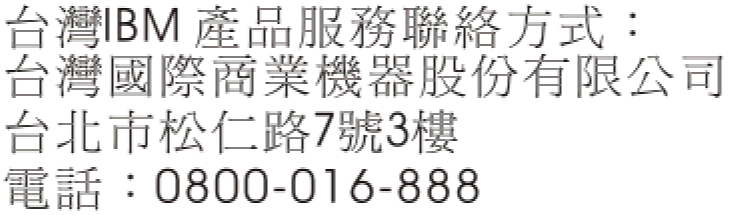

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) “. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

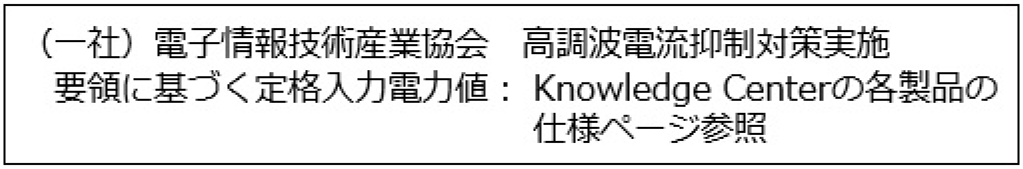

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値： Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

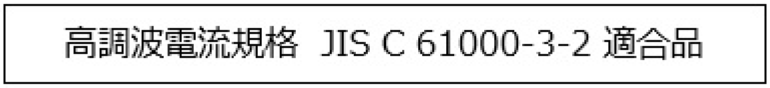

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

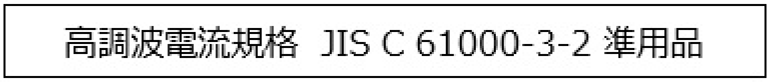

高調波電流規格 JIS C 61000-3-2 準用品

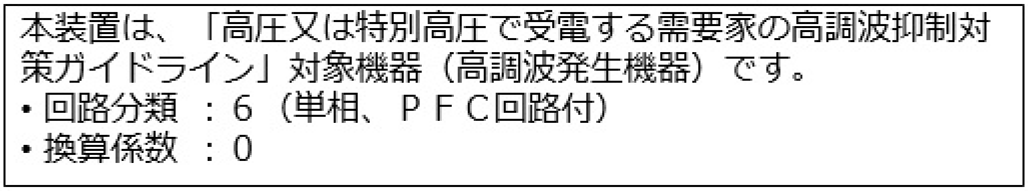

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：6（単相、ＰＦＣ回路付）
- 換算係数 ：0

This statement applies to products greater than 20 A per phase, three-phase.

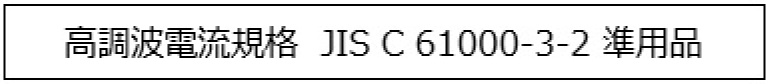

高調波電流規格 JIS C 61000-3-2 準用品

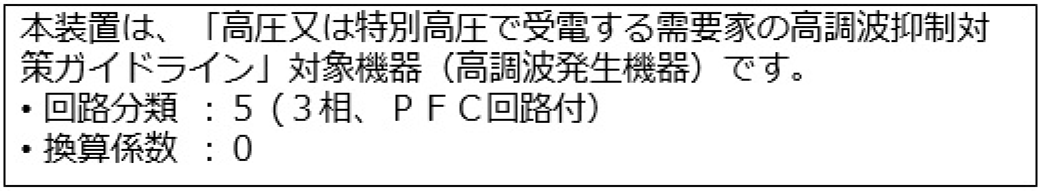

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：5（３相、ＰＦＣ回路付）
- 換算係数 ：0

Japan Voluntary Control Council for Interference (VCCI) Notice

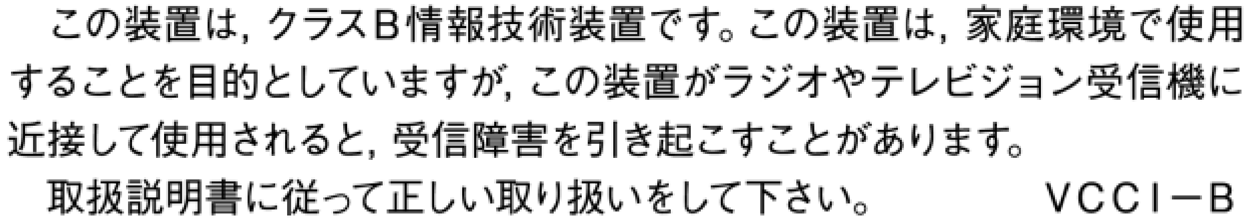

この装置は，クラスB情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。 VCCI－B

Taiwan Notice

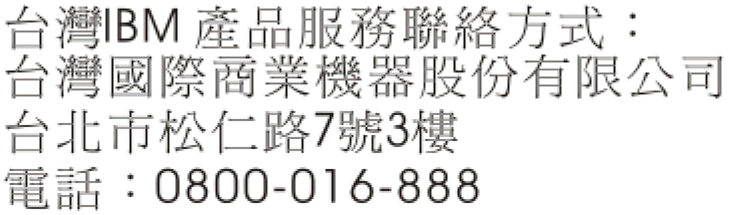

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Power Systems

Installing and configuring the Hardware Management Console

Note

Before using this information and the product it supports, read the information in "Safety notices" on page v, "Notices" on page 93, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Hardware Management Console Version 10 Release 1 Maintenance Level 1010 and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.

- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- 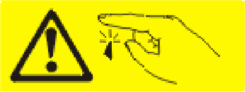 Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).

- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.

  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer’s DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

⚠ **CAUTION:**

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position (compliance ID RACK-001 or 22U (compliance ID RR001) and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.

  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U (compliance ID RACK-001 or 22U (compliance ID RR001) level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

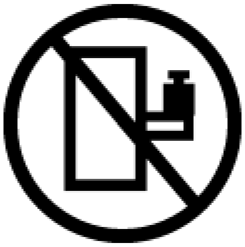

⚠ **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.

- Before extending the rack to the installation position, read the installation instructions.
- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

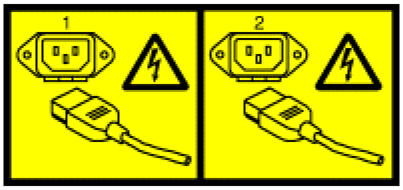

or

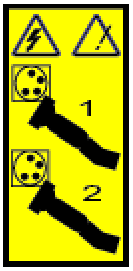

or

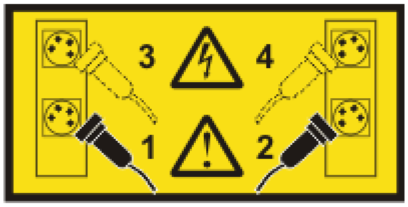

or

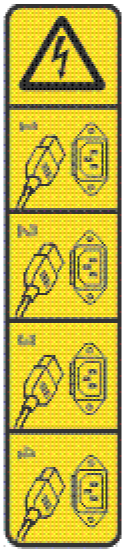

or

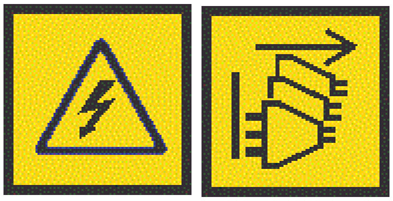

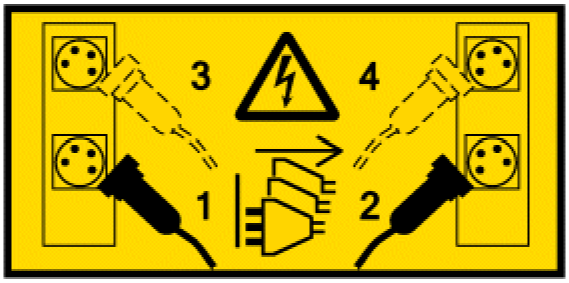

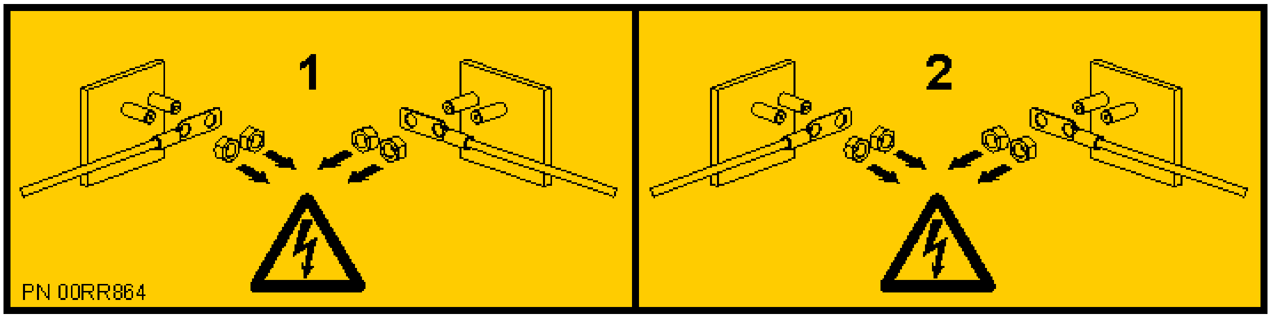

DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

CAUTION: A hot surface nearby. (L007)

(L008)

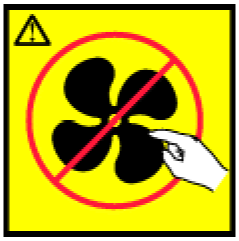

CAUTION: Hazardous moving parts nearby. (L008)

(L018)

or

CAUTION: High levels of acoustical noise are (or could be under certain circumstances) present. Use approved hearing protection and/ or provide mitigation or limit exposure. (L018)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

⚠ **CAUTION:** This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

⚠ **CAUTION:** Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

⚠ **CAUTION:** This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

⚠ **CAUTION:** Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

⚠ **CAUTION:** The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approvedapproved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)(C003a)

⚠ **CAUTION:** Regarding IBM providedprovided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.

- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four (4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.
- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.
- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intra-building ports of this equipment are suitable for connection to intra-building or unexposed wiring or cabling only. The intra-building ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The AC-powered system does not require the use of an external surge protection device (SPD).

The DC-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The DC-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Installing and configuring the Hardware Management Console

Learn how to install the Hardware Management Console (HMC) hardware, connect it to your managed system, and configure it for use. You can perform these tasks yourself, or contact a service provider to perform these tasks for you. You might be charged a fee by the service provider for this service.

What's new in Installing and configuring the HMC

Read about new or significantly changed information in the Installing and configuring the HMC topic since the previous update of the topic collection.

September 2021

- Removed the menu icons in the HMC topics.
- Updated the following topics:
  - "Installation and configuration tasks" on page 1
  - "Securing the HMC" on page 84
  - "Enhanced password policy" on page 86
  - "HMC port locations" on page 92

Installation and configuration tasks

Learn about the tasks that are associated with different HMC installation and configuration tasks.

Learn about, at a high level, the tasks you must complete when you install and configure your HMC. You can install and configure your HMC in different ways. Find the situation that best matches the task that you want to complete.

Notes:

- If you are managing Power10 processor-based systems, the HMC must be at Version 10.1.1010, or later. For more information, see "Determining your HMC machine code version and release" on page 75.
- Hardware Management Console Version 10.1.1010, or later is not supported on the HMC 7042 Machine Type. For more information about the HMC versions for your 7042 HMC, see the HMC release notes that is available in the Fix Central website.

Installing and configuring a new HMC with a new server

Learn more about the high-level tasks you must complete when you install and configure a new HMC with a new server.

Table 1. Tasks that you need to complete when you install and configure a new HMC with a new server

Task	Where to find related information
1. Gather information and complete the Preinstallation Configuration worksheet.	"Preinstallation configuration worksheet for the HMC" on page 24 "Preparing for HMC configuration" on page 23
2. Unpack the hardware.	

Table 1. Tasks that you need to complete when you install and configure a new HMC with a new server (continued)

Task	**Where to find related information**
3. Cable the HMC hardware.	"Cabling the rack-mounted 7063-CR2 HMC" on page 9
4. Power on the HMC by pressing the power button.	
5. Log in and start the HMC web application.	
6. Access the Guided setup wizard or use the HMC menus to configure the HMC.	"Configuring the HMC by using the menus " on page 31
7. Attach the server to the HMC.	

Updating and upgrading your HMC code

Learn more about the high-level tasks you must complete when you update and upgrade your HMC code.

If you have an existing HMC and want to update or upgrade your HMC code, you must complete the following high-level tasks:

Table 2. Tasks that you need to complete when you update or upgrade HMC code

Task	**Where to find related information**
1. Obtain the upgrade.	"Upgrading your HMC software" on page 80
2. View the existing HMC machine code level.	
3. Back up the managed system's profile data.	
4. Back up HMC data.	
5. Record the current HMC configuration information.	
6. Record remote command status.	
7. Save upgrade data.	
8. Upgrade the HMC software.	
9. Verify that the HMC machine code upgrade installed successfully	

Adding a second HMC to an existing installation

Learn more about the high-level tasks you must complete when you add a second HMC to your managed system.

If you have an existing HMC and managed system and want to add a second HMC to this configuration, complete the following steps:

Table 3. Tasks that you need to complete when you add a second HMC to an existing installation

Task	**Where to find related information**
1. Ensure that your HMC hardware supports HMC Version 7 code.	
2. Gather information and complete the Preinstallation Configuration worksheet.	"Preinstallation configuration worksheet for the HMC" on page 24

Table 3. Tasks that you need to complete when you add a second HMC to an existing installation (continued)

Task	Where to find related information
3. Unpack the hardware.	
4. Cable the HMC hardware.	"Cabling the rack-mounted 7063-CR2 HMC" on page 9
5. Power on the HMC by pressing the power button.	
6. Log in to the HMC.	
7. The HMC code levels must match. Change the code on one of the HMCs to match the code on the other.	"Determining your HMC machine code version and release" on page 75 "Upgrading your HMC software" on page 80
8. Access the Guided setup wizard or use the HMC menus to configure the HMC.	"Configuring the HMC by using the menus " on page 31
9. Configure this HMC for service by using the Call-Home Setup wizard.	"Configuring the HMC so that it can connect to service and support by using the call-home setup wizard" on page 69
10. Attach the server to the HMC.	

Setting up the HMC

You must set up the HMC hardware before you configure the HMC software. Learn more about setting up a desk-side HMC or a rack-mounted HMC.

Installing the IBM Power Systems HMC (7063-CR2) into a rack

Learn how to install the IBM Power® systems HMC (7063-CR2) into a rack.

You can view the online installation documentation, or you can print the PDF version of the same information. To view or print the PDF version, see Installing and configuring the Hardware Management Console.

Prerequisites for installing the rack-mounted 7063-CR2 system

Use the information to understand the prerequisites that are required for installing the system.

About this task

⚠ **CAUTION:** This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

You might need to read the following documents before you begin to install the server:

- The latest version of this document is maintained online, see Installing the 7063-CR2 into a rack (http://www.ibm.com/support/knowledgecenter/POWER9/p10hai/p10hai_install7063cr2_kickoff.htm).
- To plan your server installation, see Site and hardware planning.

Procedure

1. Ensure that you have the following items before starting your installation:
   - Size 2 Phillips screwdriver
   - Flat-head screwdriver

- T25 screwdriver
- Box cutter
- Electrostatic discharge (ESD) wrist strap
- Rack with one Electronic Industries Association (EIA) unit (1U) of space.

Notes:

- If you do not have a rack that is installed, install the rack. For instructions, see Racks and rack features (http://www.ibm.com/support/knowledgecenter/POWER10/p10hbf/p10hbf_10xx_kickoff.htm).
- The power supply ratings are 100 to 127 V ac, 9 A (x2), 200 to 240 V ac, 4.5 A (x2); 50 or 60 Hz.

2. Remove the shipping brackets on the system.
3. Continue with "Completing inventory for your system" on page 4.

Completing inventory for your system

Use this information to complete inventory for your system.

Procedure

1. Verify that you received all the boxes you ordered.
2. Unpack the server components as needed.
3. Complete a parts inventory and verify that you have received all the parts that you ordered before you install each server component.

   Note:

   Your order information is included with your product. You can also obtain order information from your marketing representative or the IBM Business Partner.

   If you have incorrect, missing, or damaged parts, consult any of the following resources:

   - Your IBM reseller.
   - IBM Rochester manufacturing automated information line at 1-800-300-8751 (United States only).
   - The Directory of worldwide contacts website (http://www.ibm.com/planetwide). Select your location to view the service and support contact information.
4. Continue with "Determining and marking the location in the rack for the 7063-CR2 system" on page 4.

Determining and marking the location in the rack for the 7063-CR2 system

You need to determine where to install the system unit into the rack.

Procedure

1. Read the Rack safety notices (http://www.ibm.com/support/knowledgecenter/POWER10/p10hbf/p10hbf_racksafety.htm).
2. Determine where to place the system unit in the rack. As you plan for installing the system unit in a rack, consider the following information:
   - Organize larger and heavier units into the lower part of the rack.
   - Plan to install system units into the lower part of the rack first.
   - Record the Electronic Industries Alliance (EIA) locations in your plan.
3. If necessary, remove the filler panels to allow access to the inside of the rack enclosure where you plan to place the unit, as shown in Figure 1 on page 5.

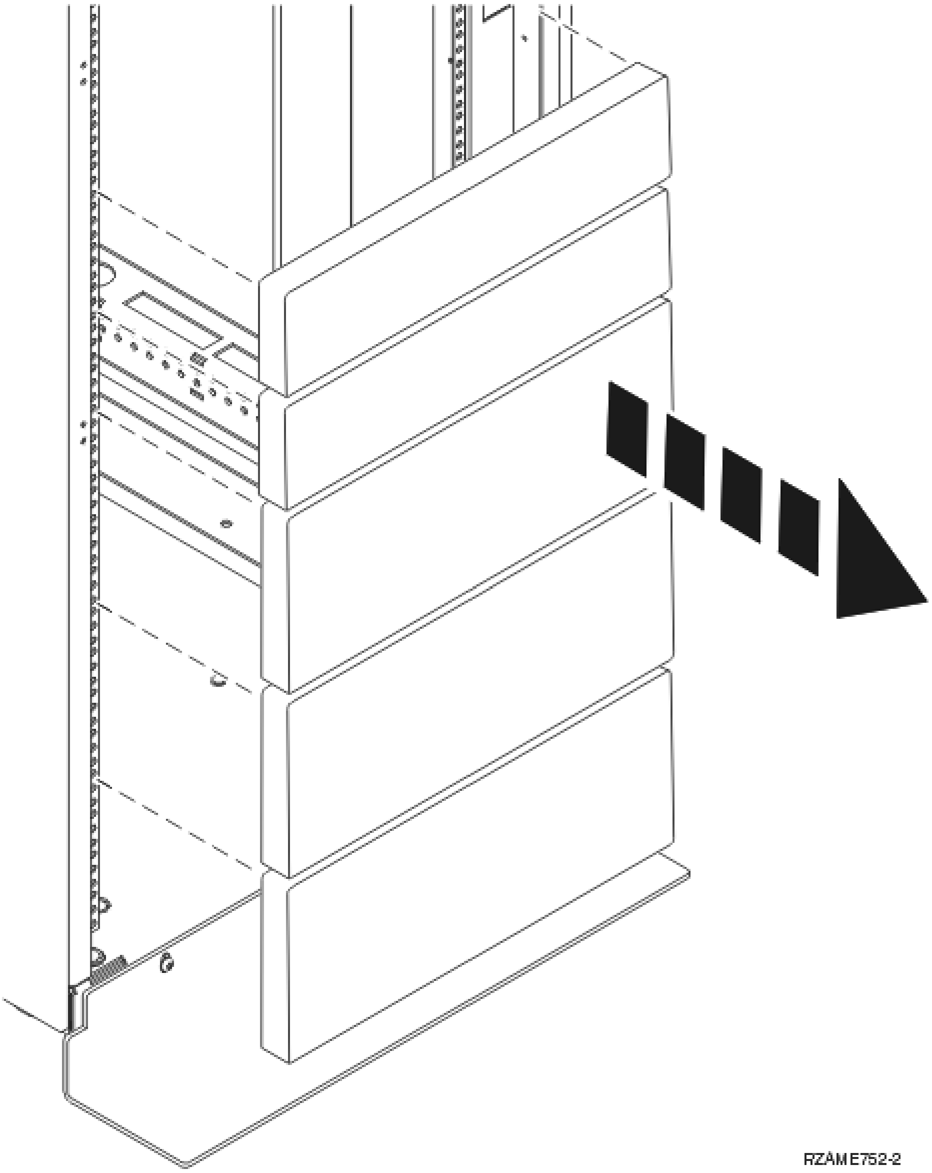

Figure 1. Removing the filler panels

4. Determine where to place the system in the rack. Record the EIA location.
5. Facing the front of the rack and working from the right side, use tape, a marker, or pencil to mark the lower hole of each EIA unit.
6. Repeat step "5" on page 5 for the corresponding holes located on the left side of the rack.
7. Go to the rear of the rack.
8. On the right side, find the EIA unit that corresponds to the bottom EIA unit marked on the front of the rack.
9. Mark the bottom EIA unit.
10. Mark the corresponding holes on the left side of the rack.
11. Continue with "Attaching the adjustable rails to the system chassis and to the rack" on page 6 to attach the adjustable rails or continue with "Attaching the fixed rails to the system chassis and to the rack" on page 7 to attach the fixed rails.

Attaching the adjustable rails to the system chassis and to the rack

You must install the rails onto the chassis and into the rack. Use this procedure to perform this task.

About this task

Attention: To avoid rail failure and potential danger to yourself and to the unit, ensure that you have the correct rails and fittings for your rack. If your rack has square support flange holes or screw-thread support flange holes, ensure that the rails and fittings match the support flange holes that are used on your rack. Do not install mismatched hardware by using washers or spacers. If you do not have the correct rails and fittings for your rack, contact your IBM reseller.If you don’t have the correct rails and fittings for your rack, contact your reseller.

Note: The 1 EIA units in racks are measured in vertical increments of 44.45 mm (1.75 in.) each. Each 44.45 mm (1.75 in.) increment is called an “EIA.” In some countries, the same increment can be referred to as a “U.”

Note: The system requires 1 EIA rack unit (1U) of space.

Ensure that you have the necessary parts to install the rails. The following parts are included with the rail kit:

- 4 - Philips 6.35 mm (0.25 in.) screws
- 2 - rack and slide bracket rail assemblies
- 2 - HMC slide brackets
- 10 - nut clips for square EIA mounting holes
- 10 - nut clips for round EIA mounting holes
- 10 - M5 hex flange screws

Procedure

1. Remove the rail pieces from the packaging and put them on a work surface.
2. Identify 1U space in the rack of the HMC.
3. To attach the slide brackets to the HMC, perform the following tasks:
   a. Identify the right slide bracket.
   b. Align the holes on the right slide bracket with the slide bracket pins located on the right side of the HMC. Ensure that all the pins are aligned with the bracket holes.
   c. Push the HMC slide bracket toward the rear side of the HMC until it is fully locked into position.
   d. Secure the right slide bracket to the right side of the HMC workstation by installing two Philips 6.35 mm (0.25 in.) screws into the screw holes.
   e. Repeat the steps “3.a” on page 6 - “3.d” on page 6 to install the left slide bracket to the left side of the HMC workstation.
4. Move to the front of the rack.
   a. On the left side, install three nut clips into the three holes on the front edge of the rack in the 1U slot that is designated for the HMC.

      Note: The rail kit includes nut clips for both square and round rack holes. Ensure that you use appropriate nut clips that match the holes in the rack.
   b. Repeat step “4.a” on page 6 on the right side of the rack.
5. Move to the rear of the rack.
   a. On the left side, install two nut clips into the upper and lower holes on the front edge of the rack in the 1U slot that is designated for the HMC.

      Note: The middle hole must remain empty.
   b. Repeat step “5.a” on page 6 on the right side of the rack.

6. To install the HMC slide rails into the rack, perform the following steps:
    a. Measure the depth of the rack. The depth must be between 558.8 mm (22 in.) and 863.6 mm (34 in.).
    b. Place the HMC slide rails on a flat surface and locate the preinstalled screws.

       Note: The slide rails have four screw holes.
    c. Loosen the preinstalled screws on the slide rails enough that the rails can be moved in and out easily.
    d. Based on the depth of the rack measured in step "6.a" on page 7, you must adjust the screws on the rails.
        i) If the depth of the rack is between 558.8 mm (22 in.) and 698.5 mm (27.5 in.), attach the screws to the first and the third holes.
        ii) If the depth of the rack is between 698.5 mm (27.5 in.) and 863.6 mm (34 in.), attach the screws to the second and the fourth holes.

        Notes:

        - The first hole is always the hole closest to the end of the slide rail. The third and fourth holes are located close together.
        - Ensure that the screws are loose enough so that the length of the slide rail can be slightly adjusted while it is being installed in the rack.
7. At the front of the rack, install the HMC slide rails into the rack by performing the following steps:
    a. Locate the left slide rail assembly.
    b. Orient the rail assembly so that the end with the closest screw hole (the first hole) goes into the rack first. Ensure that the screw heads face the inside of the rack. The open slot of the rail assembly is closest to the front of the rack.
    c. On the left side of the rack, connect the flange on the end of the slide rail to the front edge of the rack by using two M5 screws, leaving the middle hole open. Ensure that the rail assembly is left slightly loose on the front of the rack to allow for the HMC to be inserted.
8. At the rear of the rack, on the right side, pull the free end of the slide rail toward the rear and secure the flange of the slide rail to the rack by using two M5 screws, leaving the middle screw hole open.
9. Repeat the step "7" on page 7 and step "8" on page 7 to install the right slide rail assembly on the right side of the rack.
10. In front of the rack, install the HMC workstation into the rack by performing the following steps:
    a. Holding the HMC workstation level, insert the slide brackets into the HMC slide rails that you installed in the previous step. Push the HMC forward until the flanges on the front of the HMC are flush with the open screw holes on the front of the rack.
    b. Connect the HMC to the left side of the frame using one M5 screw. Repeat this step on the right side of the rack.
11. Continue with "Installing the system into the rack and connecting and routing power cables" on page 9.

Attaching the fixed rails to the system chassis and to the rack

You must install the rails onto the chassis and into the rack. Use this procedure to perform this task.

About this task

Attention: To avoid rail failure and potential danger to yourself and to the unit, ensure that you have the correct rails and fittings for your rack. If your rack has square support flange holes or screw-thread support flange holes, ensure that the rails and fittings match the support flange holes that are used on your rack. Do not install mismatched hardware by using washers or spacers.

If you do not have the correct rails and fittings for your rack, contact your IBM reseller.If you don’t have the correct rails and fittings for your rack, contact your reseller.

Note: The 1 EIA unit in racks are measured in vertical increments of 44.45 mm (1.75 in.) each. Each 44.45 mm (1.75 in.) increment is called an “EIA.” In some countries, the same increment can be referred to as a “U.”

Note: The system requires 1 EIA rack unit (1U) of space.

Ensure that you have the necessary parts to install the rails. The following parts are included with the rail kit:

- 4 - Philips 6.35 mm (0.25 in.) screws
- 2 - Inner rails
- 2 - HMC support rails
- 2 - Nut clips for square EIA mounting holes
- 2 - Nut clips for round EIA mounting holes
- 8 - M5 hex flange screws

Procedure

1. Remove the rail pieces from the packaging and put them on a work surface.
2. Identify 1U space in the rack of the HMC.
3. To attach the inner rails to the HMC, perform the following tasks:
   a. Identify the right inner rail.
   b. Align the holes on the right inner rail with the inner rail pins located on the right side of the HMC. Ensure that all the pins are aligned with the inner rail holes.
   c. Push the HMC inner rail toward the front side of the HMC until it is fully locked into position.
   d. Secure the right inner rail to the right side of the HMC workstation by installing two Philips 6.35 mm (0.25 in.) screws into the screw holes.
   e. Repeat the steps 3.a - “3.d” on page 8 to install the left inner rail to the left side of the HMC workstation.
4. Move to the front of the rack. On the left side, install one nut clips into the hole on the front edge of the rack in the 1U slot that is designated for the HMC.

   Note: The rail kit includes nut clips for both square and round rack holes. Ensure that you use appropriate nut clips that match the holes in the rack.
5. Move to the rear of the rack. On the left side, install one nut clips into the middle hole on the front edge of the rack in the 1U slot that is designated for the HMC.
6. At the front of the rack, install the HMC support rails into the rack by performing the following steps:
   a. Align the pins of the support rails above and below the nut clip you installed in the previous step..
   b. On the right side of the rack, connect the flange on the end of the support rail to the front edge of the rack by using two M5 screws into the upper and lower screw holes, leaving the middle screw hole open. Ensure that the rail assembly is left slightly loose on the front of the rack to allow for the HMC to be inserted.
7. At the rear of the rack, on the right side, pull the free end of the support rail toward the rear and secure the flange of the support rail to the rack by using two M5 screws, leaving the middle screw hole open.
8. Repeat the step “6” on page 8 and step “7” on page 8 to install the right support rail assembly on the right side of the rack.
9. In front of the rack, install the HMC workstation into the rack by performing the following steps:

a. Holding the HMC workstation level, insert the inner rails into the HMC support rails that you installed in the previous step. Push the HMC forward until the flanges on the front of the HMC are flush with the open screw holes on the front of the rack.

b. Connect the HMC to the left side of the frame using one M5 screw. Repeat this step on the right side of the rack.

Note: If present, remove the orange shipping brackets that are attached to the rear of the system and reinstall the screw back in.

10. Continue with "Installing the system into the rack and connecting and routing power cables" on page 9.

Installing the system into the rack and connecting and routing power cables

Install the system onto the rails and connect and route power cables.

About this task

⚠ **CAUTION:** This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

Procedure

1. Plug the power cords into the power supplies.

   Note: Do not connect the other end of the power cord to the power source now.

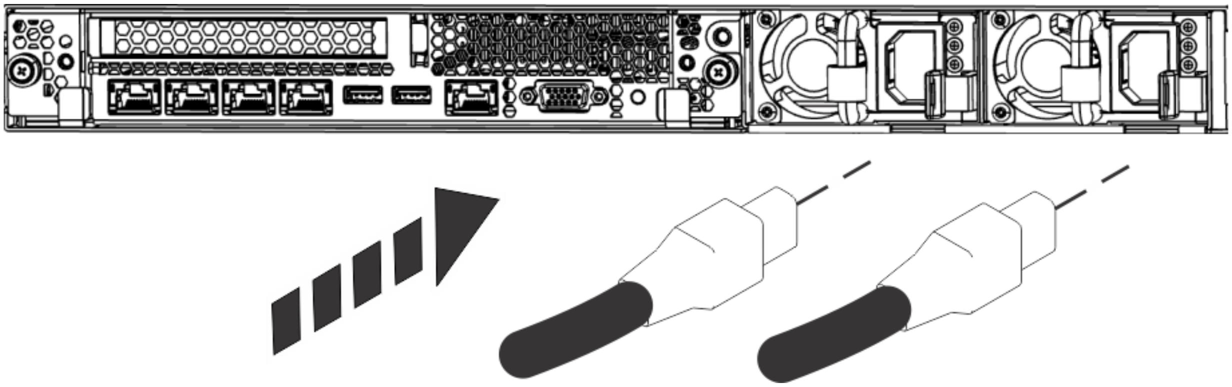

Figure 2. Plugging the power cords into the power supplies

2. Fasten the hook-and-loop fasteners to secure the power cords.
3. Continue with "Cabling the rack-mounted 7063-CR2 HMC" on page 9.

Cabling the rack-mounted 7063-CR2 HMC

Learn how to physically install your rack-mounted Hardware Management Console (HMC).

Procedure

1. Ensure that the HMC is installed into a rack and the power cords are plugged into the power supplies. For more information, see "Installing the system into the rack and connecting and routing power cables" on page 9. After you install the HMC into a rack, continue with the next step.
2. Connect the keyboard, monitor, and mouse.

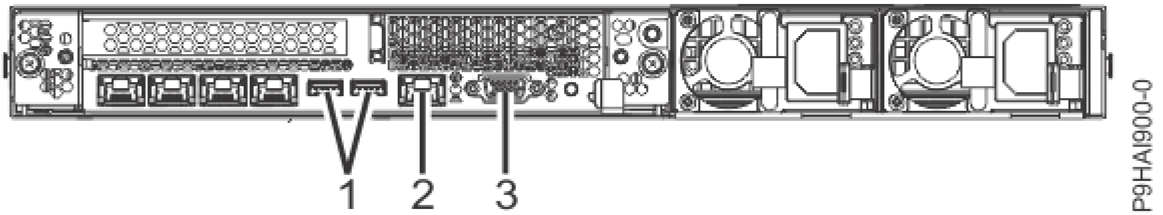

Figure 3. Rear ports

Table 4. Input and output ports	
Identifier	**Description**
1	USB 2.0 used for keyboard and mouse
2	Ethernet Intelligent Platform Management Interface (IPMI)
3	Video Graphics Array (VGA) that is used for the monitor. Only the 1024 x 768 at 60 Hz VGA setting is supported. Only up to a 3-meter cable is supported.

Note: The system has two front USB ports that you can use.

3. Connect the Ethernet Intelligent Platform Management Interface (IPMI) port to a network.

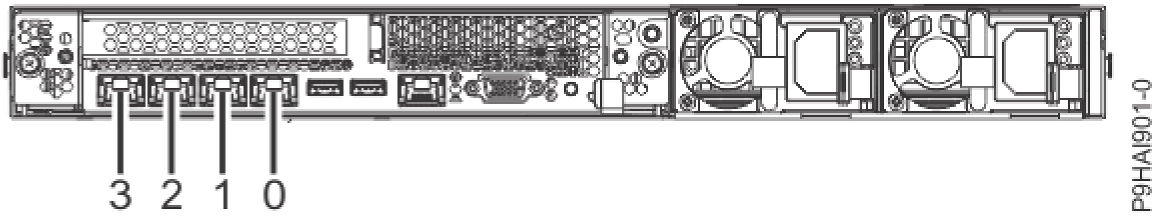

Figure 4. Ethernet ports

Table 5. Ethernet ports	
Identifier	**Description**
0	Shared Ethernet Intelligent Platform Management Interface (IPMI) and HMC Network Connection
1, 2, and 3	HMC network connection

Notes:

- When the HMC is equipped with an optional 10 Gb PCI Express (PCIe) adapter, two additional ethernet ports are available.
- This connection is required to access the baseboard management controller (BMC) on the HMC. Access to the BMC is required for service tasks and to maintain the HMC firmware. For more information, see “Types of HMC network connections” on page 15.

Warning: This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Please contact IBM for more information.

4. Connect the Ethernet cable that is intended for the connection to the managed system or systems.

Notes:

- If you are using a shared connection for IPMI and HMC, a single cable to port 0 in Figure 2 can satisfy both requirements for IPMI and HMC.
- To learn more about the HMC network connections, see "HMC network connections" on page 15.

5. If your managed system is already installed, you can verify that the Ethernet cable connection is active by observing the green status lights at both the HMC and managed system Ethernet ports as your installation progresses.
6. Plug the system power cords and the power cords for any other attached devices into the alternating current (AC) power source.
7. Verify the power status by using the power supply LEDs as indicators. For more information, see LEDs on the 7063-CR2 systemLEDs on the 7063-CR2 system.
8. Press the power button to start the system. The power-on light stops flashing and remains on, indicating that the system power is on.

Results

Next, you need to install and configure your HMC software. Continue with "Configuring the 7063-CR2 HMC" on page 11.

Configuring the 7063-CR2 HMC

Learn how to install and configure the Hardware Management Console (HMC).

Check the HMC version that is shipped with your HMC. To find out how to view the HMC machine code version and release, see Check the HMC version that is shipped with your HMC. You can download the latest HMC version that is available from the Fix Central website. Use removable media (such as a DVD or USB) to create a bootable ISO file from the HMC package (ISO image).

Note: The following table describes the predefined (default) login information for the HMC and BMC interfaces.

Table 6.

Console or Interface	Default ID	Default Password	Description
BMC (OpenBMC)	`root`	`0penBmc`	The root user ID and password are used to log in to the BMC for the first time.
HMC	`hscroot`	`abc123`	The hscroot user ID and password are used to log in to the HMC for the first time. They are case-sensitive and can only be used by a member of the super administrator role.
HMC	`root`	`passw0rd`	The root user ID and password are used by the service provider to perform maintenance procedures. They cannot be used to log in to the HMC.

Note: The following installations are shown as examples.

Installing the HMC by using USB flash drive

To install the HMC by using USB flash drive, complete the following steps for Linux® systems:

Note: For examples in different operating systems, see:

- Windows: USB flash installation media (Windows)
- Mac: USB flash installation media (macOS)

1. Download the HMC version that you want from the Fix Central website.
2. Run the following command to identify the device name of the USB drive when it is plugged in: **`lsblk`**.

   For example: **`/dev/sdb`** (where **sdb** is the name of the USB drive)
3. Run the following command to wipe the USB drive: **`wipefs --all /dev/sdX`**.

   For example: **`wipefs --all /dev/sdb`**
4. Run the following command to verify the size of the disk under the SIZE column: **`lsblk`**.

   For example: When a 16 GB USB drive shows as 14.3 GB, round it down to 14 GB for the next step "5" on page 12.
5. Run the following command to format the disk and create a partition: **`parted /dev/sdX`**

   From the parted utility, run the following three commands:

   ```
   mklabel gpt
   mkpart primary ext3 1MiB <size>GiB
   quit
   ```

   Note: size is the size of the USB drive obtained in the step "4" on page 12.

   For example:

   ```
   parted /dev/sdb
   mklabel gpt
   mkpart primary ext3 1MiB 14GiB
   quit
   ```
6. Run the following command to copy the ISO onto the partition: **`cat HMC-Recovery-ppc64le.iso > /dev/sdX1`**.

   For example: **`cat HMC-9.2.950.0-2103300827-ppc64le.iso > /dev/sdb1`**
7. Insert the USB drive, and power on the system.

   Note: The USB drive must be at least 8 GB. Certain USB drives might be too wide to fit properly into the USB port at the rear of the system. Test the fit of your USB drive before you proceed.
8. If viewing the HMC console via the BMC UI (under **Server control**) instead of the local console, the **KVM** console loads with Petitboot. The pre-Petitboot output can only be seen in the **Serial over LAN console**.
9. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **USB**.

Installing the HMC by using virtual media from the BMC

To install the HMC by using virtual media from the BMC, complete the following steps:

1. Open a supported web browser. In the address bar, enter the IP address of the BMC that you want to connect to. For example, you can use the format `https://<BMC IP>` in the address bar of the web browser.
2. From the **OpenBMC logon** window, enter the **Host** address of the BMC and the **Username** and **Password** that is assigned to you.

Note: The default user ID is `root` and the default password is `0penBmc`.

If you are using firmware level OP940.01, or later, the root password is expired by default. You must change the default password before you can access the BMC. For more information about changing the expired default password, see .

If you forgot your password, you can perform a factory reset of the system to restore the default password. To reset the system, see .

3. Click **Log in**.
4. Select **Server control**.
5. Select **Virtual Media**.
6. Click **Choose file**.
7. Locate the HMC Recovery media ISO and click **Open**.
8. Click **Start**.
9. Power on the system.
10. If viewing the HMC console via the BMC UI (under **Server control**) instead of the local console, the **KVM** console loads with Petitboot. The pre-Petitboot output can only be seen in the **Serial over LAN console**.
11. When the Petitboot menu is displayed, select the **Install Hardware Management Console** option that is located under **USB**.

Installing the HMC by using an external USB attached DVD drive

To install the HMC by using an external USB attached DVD drive, complete the following steps:

1. Download the HMC recovery version that you want from the Fix Central website.
2. Burn the HMC recovery DVD image to a DVD-R DL media as an image.
3. Power off the HMC.
4. Connect the external USB DVD drive to the HMC and insert the HMC recovery DVD.

   Note: You might need to connect the USB DVD drive to an external power source or use a USB Y cable to connect to an extra USB port to provide sufficient power to the DVD drive.
5. Power on the HMC.

   Note: The display monitor might show no signal during startup. The process might take 2 or 3 minutes before the display monitor shows any status.
6. When the Petitboot bootloader starts, navigate to stop the automatic boot.

   Note: A 10-second timeout is enforced. If no action is taken within 10 seconds, the system attempts to boot from the hard disk drive.
7. Wait until the **CD/DVD** device appears in the Petitboot menu.

   Note: This process can take up to a minute.
8. If viewing the HMC console via the BMC UI (under **Server control**) instead of the local console, the **KVM** console loads with Petitboot. The pre-Petitboot output can only be seen in the **Serial over LAN console**.
9. Select the **Install Hardware Management Console** option that is located under **CD/DVD**.

Enabling secure boot on 7063-CR2 HMC

Use the information to understand the prerequisites and steps to enable secure boot on 7063-CR2 HMC.

About this task

- To enable secure boot on the HMC, the firmware verification should be enabled and the operating system should be **enforcing** verification.

- The secure boot jumper is set to the **enabled** position on all 7063-CR2 HMC systems by default.
- Secure boot is enabled on firmware on all 7063-CR2 HMC systems that are shipped with HMC version 10.1.1010.
- OS secure boot can be enabled manually on 7063-CR2 HMC systems starting from HMC version 10.1.1010.

Note: It is required to re-enable the secure boot after the following replacement procedures:

- Removing and replacing the system backplane in the 7063-CR2.
- Removing and replacing the trusted platform module in the 7063-CR2.

To enable secure boot on 7063-CR2 HMC, complete the following steps:

Procedure

1. Update the firmware on the HMC to the following minimum levels, or later:
   - PNOR: IBM-mowgli-ibm-OP9_v2.5_4.123
   - OpenBMC: op940.hmc-11.1
2. Upgrade the HMC to version 10.1.1010.
3. Access Petitboot through **Petitboot** > **System Information** to verify whether **FW verification** is enabled by scrolling down to the section **Petitboot System Information**.

```
Petitboot System Information

  Secure & trusted boot
  FW verification : enabled
  FW measurement  : enabled
  OS verification : disabled
```

4. If **FW verification** is disabled, complete the following steps:
   - Shut down the HMC.
   - Remove the system backplane. For more information, see Removing the system backplane from the 7063-CR2 system.
   - Flip the secure boot jumper from debug or insecure pins to secure pins.
5. If **OS verification** is disabled, complete the following steps:
   - Boot the HMC.
   - Copy the **PK.auth** and **db.auth** files from the HMC file system to a remote system by using the **sendfile** command as shown in the :

   ```
   sendfile-f /opt/hsc/data/secureboot/PK.auth-h <ip> -d <dir> -u <username> -n PK.auth -s
   sendfile-f /opt/hsc/data/secureboot/db.auth-h <ip> -d <dir> -u <username> -n db.auth–s
   ```

   Note: **`<ip>`** is the IP address of the remote system and **`<dir>`** is the directory in the remote system where the files are stored.
6. On the remote system, copy the **PK.auth** and **db.auth** files to a USB drive.

   Note: Alternatively, you can create an ISO file that contains both the files. This ISO file can then be mounted remotely via the Virtual Media feature of the BMC. Another method to create an ISO file from a directory is to use the **`mkisofs`** command in Linux as follows:

   ```
   mkisofs -o secureboot.iso <dir with auth files>
   ```
7. Boot the HMC to **Petitboot**, and then select **Exit to Shell**.
   - If the files were copied to a USB drive, insert the USB drive.
   - If the files were combined into an ISO file to be remotely mounted via virtual media, start the Virtual Media session.
8. Run the **`mount`** command to verify the automatic mount point.

- An example of automatic mount location for USB drive is:

  ```
  /var/petitboot/mnt/dev/sda1
  ```

  Note: Adjust the mount location as necessary for your drive.

- An example of automatic mount location for ISO file over Virtual Media is:

  ```
  /var/petitboot/mnt/dev/sdb
  ```

9. Change directory to the location of the keys in the media.

   For example, when the mount location is **`/var/petitboot/mnt/dev/sda1`**, you can run the command as following:

   ```
   cd /var/petitboot/mnt/dev/sda1
   ```

10. Write the contents of the **PK** key and **db** files to the system firmware by running the following commands:

    ```
    cat PK.auth > /sys/firmware/secvar/vars/PK/update
    cat db.auth > /sys/firmware/secvar/vars/db/update
    ```

11. Reboot the system by running the **`reboot`** command.
12. Stop the system at **Petitboot** again, and then select **System Information** to confirm that the **Secure & trusted boot** section reflects the following information:

    ```
    Secure & trusted boot
      FW verification : enabled
      FW measurement  : enabled
      OS verification : enforcing
    ```

13. Exit the shell. On the main **Petitboot** menu, select **Hardware Management Console**.
14. Verify that the HMC reports secure boot as enabled by running the command **`lshmc --boot`**. If secure boot is enabled, the following message is displayed:

    ```
    secure_boot=1
    ```

Configuring the HMC

Learn how to set up your network connections, configure your HMC, complete postconfiguration steps, and upgrade and update your HMC.

Choosing network settings on the HMC

Learn about the network settings that you can use on the Hardware Management Console (HMC).

HMC network connections

Learn how the Hardware Management Console HMC can be used in a network.

You can use different types of network connections to connect your HMC to managed systems. For more information about using the HMC on a network, see the following information:

Types of HMC network connections

Learn how to use the HMC remote management and service functions by using your network.

The HMC supports the following types of logical communications:

HMC to managed system

Used to perform most of the hardware management functions, in which HMC issues control function requests through the service processor of the managed system. The connection between the HMC and the service processor is sometimes referred to as the *service network*. This connection is required for managed system management.

HMC to logical partition
Used to collect platform-related information (hardware error events, hardware inventory) from the operating systems that are running on logical partitions, and to coordinate certain platform activities (dynamic LPAR, concurrent repair) with those operating systems. If you want to use service and error notification features, you must create this connection.

HMC to BMC

Note: The baseboard management controller (BMC) connection is applicable only to HMC model 7063-CR1.

Used to perform service and maintenance tasks. The BMC connection is used to load and maintain the HMC firmware on the system. This connection is required for access to the BMC on the HMC.

HMC to remote users
Provides remote users with access to HMC functions. Remote users can access the HMC in the following ways:

- By using the web browser to access all the HMC GUI functions remotely.
- By using Secure Socket Shell (SSH) to access the HMC command line functions remotely.
- By using a virtual terminal server for remote access to virtual logical partition consoles.

HMC to service and support
Used to transmit data, such as hardware error reports, inventory data, and microcode updates, to and from your service provider. You can use this communications path to make automatic service calls.

Your HMC can support up to four separate physical Ethernet interfaces, depending on the model. The stand-alone version of the HMC supports only three HMC interfaces, by using one integrated Ethernet adapter and up to two plug-in adapters. Use each of these interfaces in the following ways:

- One network interface can be used exclusively for HMC-to-managed system communications, which means that only the HMC and service processors of the managed systems would be on that network. One or more network interfaces can be used exclusively for HMC-to-managed system communications, which means that only the HMC and service processors of the managed systems are on that network. Even though the network interfaces into the service processors are encrypted for the Secure Sockets Layer (SSL) Protocol and password-protected, having a separate dedicated network can provide a higher level of security for these interfaces.
- An open network interface would typically be used for the network connection between the HMC and the logical partitions on the managed systems, for the HMC-to-logical partition communications. You can also use this open network interface to manage the HMC remotely.
- Optionally, you can use a third interface to connect to logical partitions and manage the HMC remotely. This interface can also be used as a separate HMC connection to different groups of logical partitions. For example, you might want to have an administrative LAN that is separate from the LAN on which all the usual business transactions are running. Remote administrators can access the HMC and other managed units by using this method. Sometimes the logical partitions are in different Network security domains, perhaps behind a firewall, and you might want to have different HMC network connections into each of those two domains.

Web browser requirements for HMC

The Hardware Management Console (HMC) version 9.1.0 is supported by Google Chrome version 57, Microsoft Internet Explorer (IE) version 11.0, Mozilla Firefox versions 45 and 52 Extended Support Release (ESR), and Safari version 10.1.

If your browser is configured to use an Internet proxy, a local IP addresses should be included in the exception list. Consult your network administrator for more information on the exception list. If you still need to use the proxy to get to the HMC, enable Use HTTP 1.1 through proxy connections under the Advanced tab in your Internet Options window.

Session cookies need to be enabled in order for ASMI to work when connected to HMC remotely. The asm proxy code saves session information and uses it. Follow the steps to enable the session cookies.

Enabling session cookies in Internet Explorer.

1. Select Tools and Click Internet Options
2. Select Privacy and Click Advanced
3. Ensure that the Always allow session cookies is checked. If not, select the Override automatic cookie handling and select Always allow session cookies.
4. Select Prompt under First-party Cookies and Third-party Cookies
5. Click OK.

Enabling session cookies in Firefox.

1. Select Tools and click Options
2. Click Cookies
3. Select Allow sites to set cookies.
4. Select Exceptions and add HMC.
5. Click OK.

Private and open networks in the HMC environment
The Hardware Management Console (HMC) can be configured to use open and private networks. Private networks allow the use of a selected range of nonroutable IP addresses. A *public*, or "open" network describes a network connection between the HMC to any logical partitions and to other systems on your regular network.

Private networks

The only devices on the HMC private network are the HMC itself and each of the managed systems to which that HMC is connected. The HMC is connected to each managed system's Flexible Service Processor (FSP).

On most systems, the FSP provides two Ethernet ports that are labeled **HMC1** and **HMC2**. You to connect up to two HMCs.

Some systems have a dual-FSP option. In this situation, the second FSP acts as a redundant backup. The basic setup requirements for a system with two FSPs are essentially the same as a system without a second FSP. The HMC must be connected to each FSP, so more network hardware is required (for example, a LAN switch or hub) when there is more than one FSP or multiple managed systems.

Note: Each FSP port on the managed system must be connected to only one HMC.

Public networks

The open network can be connected to a firewall or router for connecting to the internet. Connecting to the internet allows the HMC to call home when any hardware errors need to be reported.

The HMC itself provides its own firewall on each of its network interfaces. A basic firewall is automatically configured when you run the HMC Guided Setup wizard, but you customize your firewall settings after the initial HMC installation and configuration.

HMC as a DHCP server
You can use the Hardware Management Console (HMC) as a Dynamic Host Configuration Protocol (DHCP) server.

If you want to configure the first network interface as a private network, you can select from a range of IP addresses for the DHCP server to assign to its clients. The selectable address ranges include segments from the standard nonroutable IP address ranges.

In addition to these standard ranges, a special range of IP addresses is reserved for IP addresses. This special range can be used to avoid conflicts in cases where the HMC attached open networks are using one of the nonroutable address ranges. Based on the range that is selected, the HMC network interface on the private network is automatically assigned the first IP address of that range, and the service processors are then assigned addresses from the rest of the range.

The DHCP server in the HMC uses automatic allocation, which means that each unique service processor Ethernet interface is reassigned the same IP address each time it is started. Each Ethernet interface has a unique identifier that is based on a built-in Media Access Control (MAC) address, which allows the DHCP server to reassign the same IP parameters. You can configure both **eth0** and **eth1** HMC ports to serve DHCP addresses.You can configure both **eth0** and **eth1** HMC ports to serve DHCP addresses.

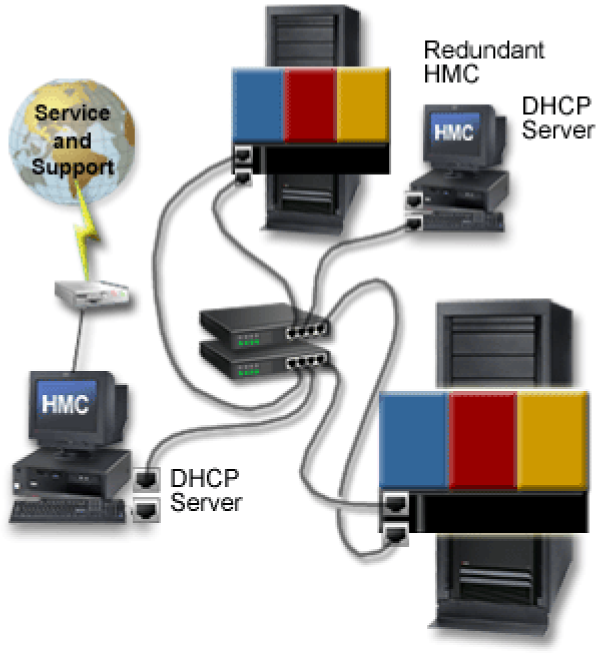

Figure 5. Private network with one HMC as a DHCP server

Note: If you are using IPv6, the discovery process must be done manually. For IPv6, automatic discovery is not available.

For more information about how to configure the HMC as a DHCP server, see "Configuring the HMC as a DHCP server" on page 38.

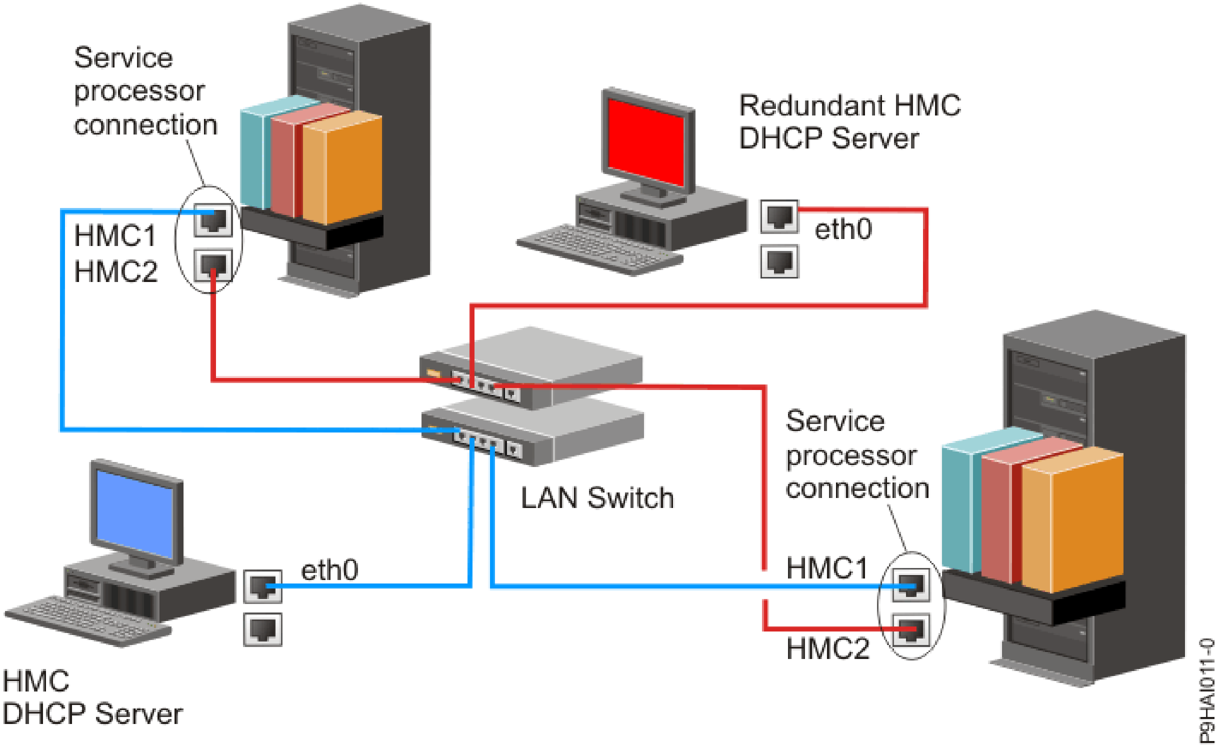
Service processor connection
Redundant HMC DHCP Server
HMC1
HMC2
eth0
LAN Switch
Service processor connection
eth0
HMC1
HMC2
HMC DHCP Server
P9HAI011-0

This figure shows a redundant HMC environment with two managed systems. The first HMC is connected to the first port on each FSP, and the redundant HMC is connected to the second port on each HMC. Each HMC is configured as a DHCP server, by using a different range of IP addresses. The connections are on separate private networks. As such, it is important to ensure that no FSP port is connected to more than one HMC.

Each managed system's FSP port that is connected to an HMC requires a unique IP address. To ensure that each FSP has a unique IP address, use the HMC's built-in DHCP server capability. When the FSP detects the active network link, it issues a broadcast request to locate a DHCP server. When correctly configured, the HMC responds to that request by allocating one of a selected range of addresses.

If you have multiple FSPs, you must have your own LAN switch or hub for the HMC to FSP private network. Alternately, this private segment can exist as several ports in a private *virtual LAN* (VLAN) on a larger managed switch. If you have multiple private VLANs, you must ensure that they are isolated and without any crossover traffic.

If you have more than one HMC, you must also connect each HMC to the logical partitions, and to each other, on the same open network.

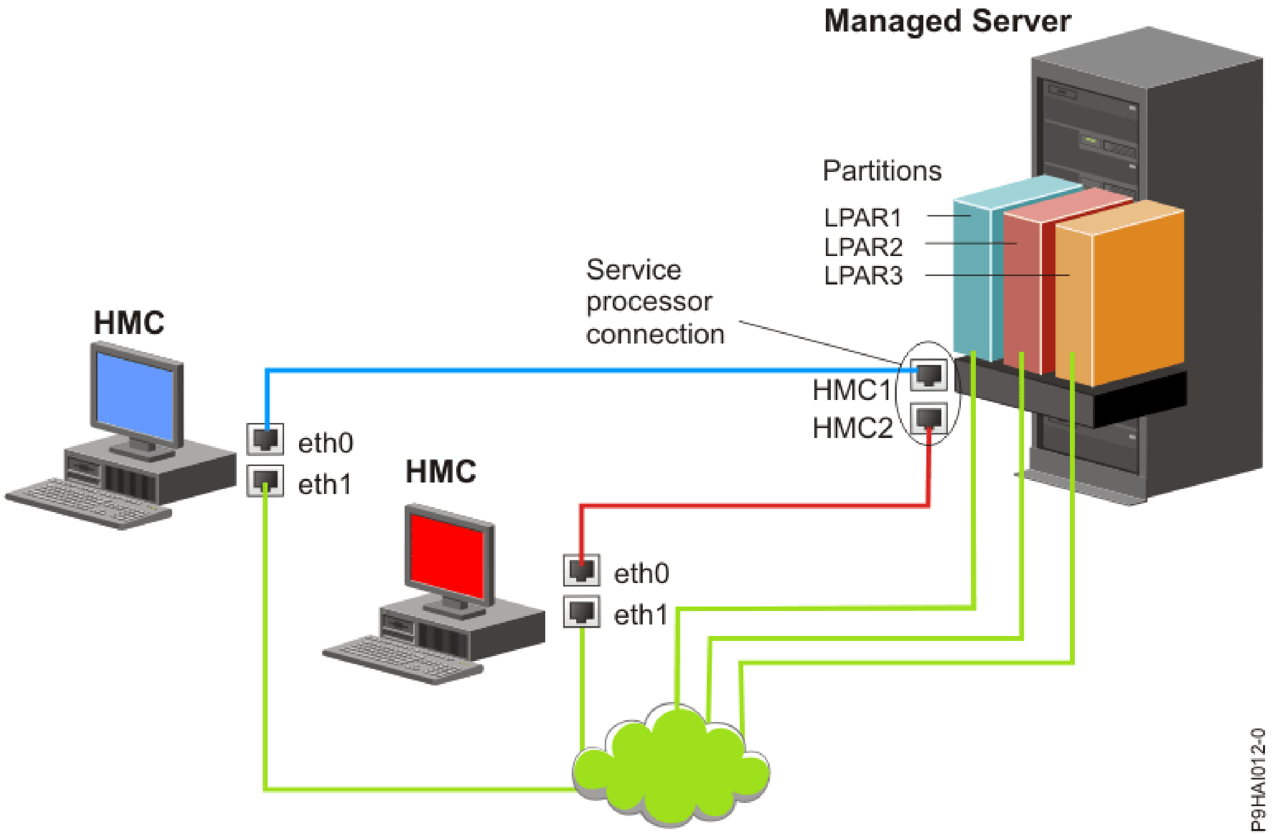

This figure shows two HMCs connected to a single managed server on the private network and to three logical partitions on the public network. You can have an extra Ethernet adapter for the HMC to have three network interfaces. You can use this third network as a management network or connect it to the CSM (Cluster Systems Manager) Management Server.

Deciding which connectivity method to use for the call-home server

Learn more about the connectivity options you have when you use the call-home server.

You can configure the Hardware Management Console (HMC) to send hardware service-related information to IBM by using a LAN-based internet connection, or a dial-up connection over a modem.

You have two communication choices when you configure the LAN-based internet connection. The first choice is to use standard Secure Sockets Layer (SSL). The SSL communication can be enabled to connect to the internet through your proxy server. SSL connectivity is more likely to be compliant with corporate security guidelines.

Note: If your open network interface connection uses only Internet Protocol Version 6 (IPv6), you cannot use internet VPN to connect to support. For more information about the protocols that are used, see "Choosing an Internet Protocol" on page 22.

The advantages to using an internet connection can include:

- Faster transmission speed
- Reduced customer expense (for example, the cost of a dedicated analog telephone line)
- Greater reliability

The following security characteristics are in effect, regardless of the connectivity method chosen:

- Remote Support Facility requests are always initiated from the HMC to IBM. An inbound connection is never initiated from the IBM Service Support System.
- All data that is transferred between the HMC and the IBM Service Support System are encrypted by using a high-grade encryption. Depending upon the connectivity method that is chosen, it is encrypted by using either SSL or IPSec Encapsulating Security Payload (ESP).

- When you initialize the encrypted connection, the HMC authenticates the target destination as the IBM Service Support System.

Data sent to the IBM Service Support System consists solely of information about hardware problems and configuration. No application or customer data is transmitted to IBM.

Using an indirect internet connection with a proxy server

If your installation requires the HMC to be on a private network, you might be able to connect indirectly to the internet by using an SSL proxy, which can forward requests to the internet. One of the other potential advantages of using an SSL proxy is that the proxy can support logging and audit facilities.

To forward SSL sockets, the proxy server must support the basic proxy header functions (as described in RFC 2616) and the CONNECT method. Optionally, basic proxy authentication (RFC 2617) can be configured so that the HMC authenticates before you attempt to forward sockets through the proxy server.

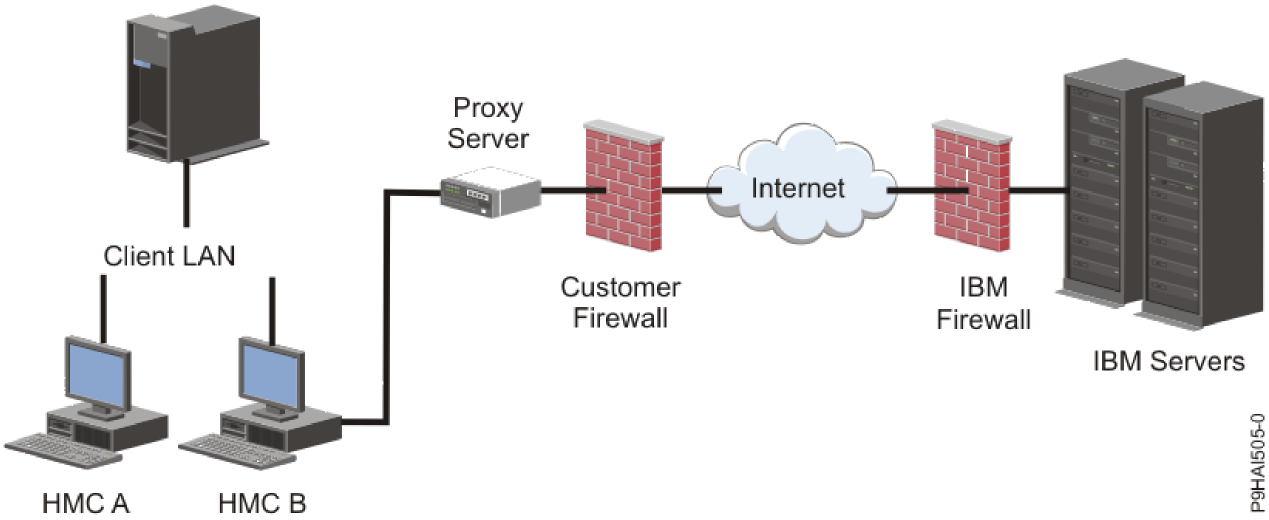

For the HMC to communicate successfully, the client's proxy server must allow connections to port 443. You can configure your proxy server to limit the specific IP addresses to which the HMC can connect. See "Internet SSL address lists" on page 22 for a list of IP addresses.

Using a direct internet SSL connection

If your HMC can be connected to the internet, and the external firewall can be set up to allow established TCP packets to flow outbound to the destinations described in "Internet SSL address lists" on page 22, you can use a direct internet connection.

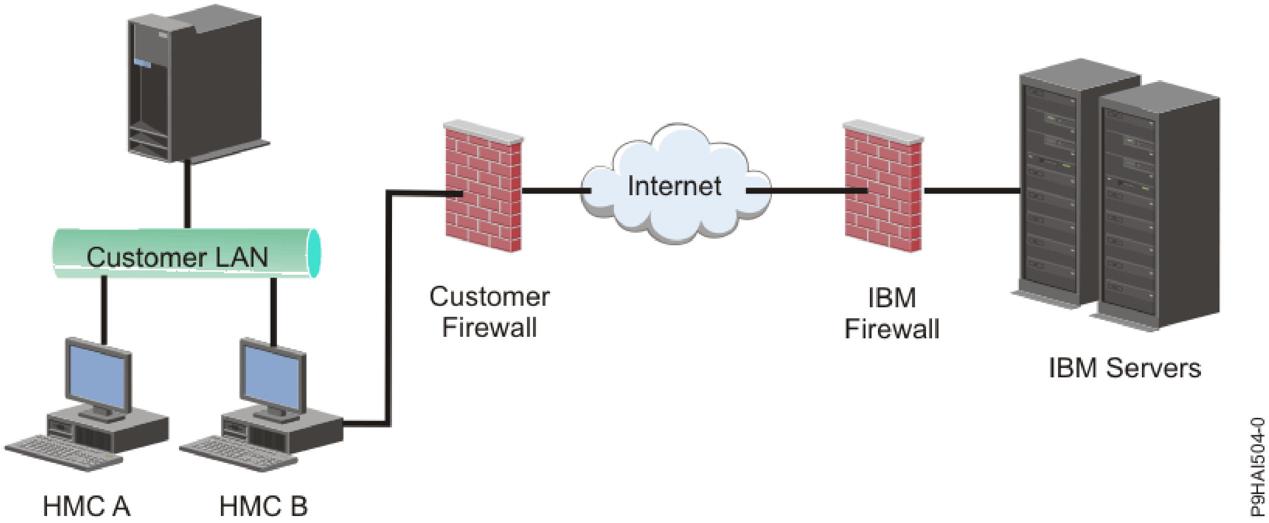

Using internet SSL to connect to remote support

All the communications are handled through TCP sockets that are initiated by the Hardware Management Console (HMC) and use a high-grade SSL to encrypt the data that is transmitted. The destination TCP/IP addresses are published (see "Internet SSL address lists" on page 22) so that external firewalls can be configured to allow these connections.

Note: The standard HTTPS port 443 is used for all communications.

The HMC can be enabled to connect directly to the internet or to connect indirectly from a proxy server that is provided by the customer. The decision about which approach is best for your installation depends on the security and networking requirements of your enterprise. The HMC (directly or through the SSL proxy) uses the following addresses when it is configured to use internet SSL connectivity.

Choosing an Internet Protocol

Determine the IP address version that is used when the Hardware Management Console (HMC) connects to your service provider.

Most users use Internet Protocol Version 4 (IPv4) to connect to a service provider. IPv4 addresses appear in the format that represents the 4 bytes of the IPv4 address, which is separated by periods (for example, 9.60.12.123) to access the internet. You can also use Internet Protocol Version 6 (IPv6) to connect to a service provider. IPv6 is often used by network administrators to ensure unique address space. If you are unsure of the Internet Protocol used by your installation, contact your network administrator. For more information about using each version, see "Setting the IPv4 address" on page 64 and "Setting the IPv6 address" on page 64.

Internet SSL address lists

Learn about the addresses that the Hardware Management Console (HMC) uses when the HMC is using internet SSL connectivity.

The HMC uses the following IPv4 addresses to contact IBM service and support when it is configured to use internet SSL connectivity.

The following IPv4 addresses are for all locations:

- 129.42.26.224
- 129.42.42.224
- 129.42.50.224
- 129.42.56.216
- 129.42.58.216
- 129.42.60.216
- 170.225.15.41

The following IPv4 addresses are for the Americas:

- 129.42.160.48
- 129.42.160.49
- 207.25.252.197
- 207.25.252.200
- 207.25.252.204

The following IPv4 addresses are for all locations other than the Americas:

- 129.42.160.48
- 129.42.160.50
- 207.25.252.197
- 207.25.252.200
- 207.25.252.205

Note: When you configure a firewall to allow an HMC to connect to these servers, only the IP addresses specific to the geographic region are required.

The HMC uses the following IPv6 addresses to contact IBM service and support when it is configured to use internet SSL connectivity:

- 2620:0:6C0:1::1000
- 2620:0:6C2:1::1000
- 2620:0:6C4:1::1000

Using multiple call-home servers

Learn about what you need to know when you decide to use more than one call-home server.

To avoid a single point of failure, configure the Hardware Management Console (HMC) to use multiple call-home servers. The first available call-home server attempts to handle each service event. If the connection or transmission fails with this call-home server, the service request is retried by using the other available call-home servers until one is successful or all servers are tried.

The connected HMC that is identified by the problem analysis to be the primary analyzing console for a given managed system that reports the problem. This primary console also replicates the problem report to any secondary HMC. This secondary HMC must be recognized on the network by the primary HMC. A secondary HMC is recognized by the primary HMC as an extra call-home server when:

- The primary HMC is configured to use "discovered" call-home servers and the call-home server is either on the same subnet as the primary HMC or it manages the same system.
- The call-home server is manually added to the list of call-home server consoles available for outbound connectivity.

Preparing for HMC configuration

Learn about the required configuration settings that you need to know before you begin the configuration steps.

To configure the HMC, you must understand the related concepts, make decisions, and prepare information.

Learn about the information that you need to connect your HMC to the following locations:

- Service processors in your managed systems
- Logical partitions on those managed systems
- Remote workstations
- IBM Service to implement “call-home” functions

To prepare for HMC configuration, complete the following steps:

1. Obtain and install the latest level of the HMC code version you want to install.
2. Determine the physical location of the HMC in relation to the servers it manages. If the HMC is more than 25 feet from its managed system, you must provide web browser access to the HMC from the managed system's location so that service personnel can access the HMC.
3. Identify the servers that the HMC manages.
4. Determine whether you use a private or an open network to manage servers. If you decide to use a private network, use DHCP, unless you are using a Cluster Systems Management (CSM) configuration. CSM does not support IPv6. To access CSM, you must have two networks. For more information about CSM, see the documentation that was provided with that feature. For more information about private and open networks, see “Selecting a private or open network” on page 38.
5. If you use an open network to manage an FSP, you must set the FSP's address manually through the Advanced System Management Interface menus. A private, non-routable network is recommended.
6. If you have two HMCs, designate a primary and secondary HMC. The primary HMC needs be physically closer to the system, and must be the HMC that is configured to call home.

7. Determine the network settings that you need to connect the HMC to remote workstations, logical partitions, and network devices.
8. Define how the HMC calls home. Call home options include either over an outbound-only Secure Socket Layer (SSL) internet connection, a modem, or a Virtual Private Network (VPN) connection.
9. Determine the HMC users that you create and their passwords, as well which roles they are given. You must assign the **hscroot** and **hscpe** users a password.
10. Document the following company contact information that is needed when you configure call home:
    - Company name
    - Administrator contact
    - Email address
    - Telephone numbers
    - Fax numbers
    - The street address of the HMC's physical location
11. If you plan to use email to notify operators or systems administrators when information is sent to IBM Service through call-home, identify the Simple Mail Transfer Protocol (SMTP) server and the email addresses you use.
12. You must define the following passwords:
    - The access password that is used to authenticate the HMC to the FSP.
    - The ASMI password that is used for the **admin** user.
    - The ASMI password that is used for the **general** user.

    Create the passwords when you connect from the HMC to a new server for the first time. If the HMC is a redundant or second HMC, obtain the HMC User password and be prepared to enter it when you connect the first time to the managed server's FSP.

When you complete these preparation steps, complete the "Preinstallation configuration worksheet for the HMC" on page 24.

Preinstallation configuration worksheet for the HMC

Use this worksheet to have the installation information you need ready for the installation.

Improved password policy for HMC

You must set a new password on the first use for newly manufactured systems with HMC version 9.940.0, or later, and after a factory reset of the system. This policy change helps to enforce that the HMC is not left in a state with a well-known password.

With HMC Version 9.940.0, and later, the `hscroot` password is expired and must be changed before you can access the functions of the HMC. For more information on how to change the password, see https://www.ibm.com/support/knowledgecenter/POWER10/p10eh6/p10eh6_useridsandpassword.htm. However, if you are upgrading from a previous HMC level or an operational installation, you do not have to change the password.

Network Settings

LAN Interface: Choose the available adapters (such as eth0, eth1) that is used by this HMC to connect to managed systems, logical partitions, service and support, and remote users. For more information, see "HMC network connections" on page 15. Connectivity from the HMC can either be on a private or open network.

Ethernet Adapter Speed and Duplex

Enter the wanted Ethernet adapter speed and duplex mode. The autodetection option determines which option is optimal if you are not sure which speed and duplex would produce optimum results for your hardware. Default = Autodetection Media speed specifies the speed in duplex mode of an

Ethernet adapter. Select Autodetection unless you need to specify a fixed media speed. Any device that is connected to the FSP (switches/HMC), must be set to Auto (Speed) / Auto (Duplex) mode, as it is the default FSP setting and cannot be changed.

Table 7. Ethernet Adapter Speed and Duplex

Characteristic	**eth0**	**eth1**	**eth2**	**eth3**
Select speed and duplex mode				
Media speed (Autodetection, 10/100/1000 Full/ Half Duplex)				

For more information about private and open networks, see "Private and open networks in the HMC environment" on page 17.

Table 8. Private or Open network

Characteristic	**eth0**	**eth1**	**eth2**	**eth3**
Specify **Private** or **Open** network for each adapter.				

Dynamic Host Configuration Protocol (DHCP) provides an automated method for dynamic client configuration. You can specify this HMC as a DHCP server. If this is the first or only HMC on the private network, enable the HMC as a DHCP server. When you enable the HMC as a DHCP server, the managed systems on the network are automatically configured and discovered by the HMC.

For Ethernet adapters specified as Private networks, complete the following table:

Table 9. DHCP server

Characteristics	**eth0**	**eth1**
Do you want to specify this HMC as a DHCP server? (yes/no)		
If yes, record the IP address range you want to use.		

If you are using the 7063-CR2 HMC, you must connect the Ethernet **IPMI** port to a network to access the baseboard management controller (BMC) on the HMC. For more information, see "Configure BMC connectivity (7063-CR2)" on page 63. Complete the following table for your BMC connection.

Table 10. BMC connection

Characteristics	**IPMI**
Do you want to configure this connection through DHCP mode? (yes/no)	
If no, list the specified static addresses below:	
IP address:	
Subnet mask:	
Gateway:	

For Ethernet adapters specified as *open* networks, complete the following tables. For more information about the different Internet Protocol versions, see "Configuring the HMC network types" on page 34.

Using IPv6

If you are using IPv6, talk to your network administrator and decide how you want to obtain IP addresses. Then, complete the following tables:

Table 11. IPv6 (static)

Characteristic	eth0	eth1	eth2	eth3
Are you using a statically assigned IP address? If yes, record that address here.				

Table 12. IPv6 (DHCP server)

Characteristic	eth0	eth1	eth2	eth3
Are you getting IP addresses from a DHCP server? (Yes/No)				

Table 13. IPv6 (IPv6 router)

Characteristic	eth0	eth1	eth2	eth3
Are you getting IP addresses from an IPv6 router?				

For more information about setting IPv6 addresses, see "Setting the IPv6 address" on page 64. For more information about using only IPv6 addresses, see "Using only IPv6 addresses" on page 64.

Using IPv4

Complete the following tables for Ethernet adapters that are specified as open networks by using IPv4.

Table 14. IPv4

Characteristics	eth0	eth1	eth2	eth3
Do you want to obtain an IP address automatically? (yes/no)				
If no, list the specified address below:				
TCP/IP Interface Address:				
TCP/IP Interface Network Mask:				
Firewall Settings:				
Would you like to configure HMC firewall settings? (yes/no)				

Table 14. IPv4 (continued)

Characteristics	**eth0**	**eth1**	**eth2**	**eth3**
If yes, list the applications and IP addresses that must be allowed through the firewall:				

TCP/IP information

A unique TCP/IP address is required for each node, both for Support Element (SE) and Hardware Management Console (HMC). The assigned network mask is used to generate a unique address, by default, for the local private LAN. If the nodes are connected into a larger network with an administered TCP/IP address, you can specify the TCP/IP address to be used. The default is generated by the system.

Firewall settings

HMC firewall settings create a security barrier that allows or denies access to specific network applications on the HMC. You can specify these control settings individually for each physical network interface, enabling you control over which HMC network applications can be accessed on each network.

If you configure at least one adapter as an Open network adapter, you must provide the following additional information to enable your HMC to access the LAN:

Table 15. Open network adapter

Local host information	
HMC host name:	
Domain name:	
Description of HMC:	
Gateway information	
Gateway Address: (nnn.nnn.nnn.nnn)	
Gateway device:	
DNS enablement	
Do you want to use DNS? (yes/no)	
If "yes", specify DNS Server Search Order below:	
1.	
2.	
Domain suffix search order:	
1.	
2.	

Local Host information

To identify your Hardware Management Console (HMC) to the network, enter the HMC's host name and domain name. Unless you are using only short host names on your network, enter a fully qualified host name. Domain name example: name.yourcompany.com

Gateway information

To define a default gateway, complete the TCP/IP address to be used for routing IP packets. The gateway address informs each computer or network device when to send data if the target station is not on the same subnet as the source.

DNS Enablement

The Domain Name System (DNS) is used to provide a standard naming convention for locating IP-based computers. By defining DNS servers, you can use host names to identify servers and Hardware Management Consoles (HMCs) rather than IP addresses.

DNS Server Search Order

Enter the IP addresses of DNS servers to be searched for mapping the host names and IP addresses. This search order is available only when DNS is enabled.

Domain Suffix Search Order

Enter the domain suffixes you are using. The HMC uses domain suffixes to append to unqualified names for DNS searches. Suffixes are searched in the order in which they are listed. This search order is available only when DNS is enabled.

Email notification

List email contact information if you want to be notified by email when hardware problem events occur on your system.

Table 16. Email notification

Characteristics	**Entry field**
Email Addresses:	
SMTP server:	
Port:	
Errors to be notified:	
Only call-home problem events	
All problem events	

SMTP server

Type the simple mail transfer Protocol (SMTP) address of the server to be notified of a system event. An example of an SMTP server name is `relay.us.ibm.com`.

SMTP is the Protocol that is used to send email. When you use SMTP, a client sends a message and communicates with the SMTP server by using the SMTP Protocol.

If you do not know the SMTP address of your server or are not sure, contact your network administrator.

Port

Type the port number of the server to be notified of a system event, or use the default port.

Email addresses to be notified

Enter configured email addresses to be notified when a system event occurs.

- Select **Only call-home problem events** to receive notification only when events occur that create a call-home function.
- Select **All problem events** to receive notification when any events occur.

Service contact information

Table 17. Service contact information

Characteristics	Entry field
Company name	
Administrator name	
Email address	
Phone number	
Alternative phone number	
Fax number	
Alternative phone number	
Street address	
Street address 2	
City or locality	
State	
Postal code	
Country or region	
Location of HMC (if same as above administrator address, specify “same”):	
Street address	
Street address 2	
City or locality	
State	
Postal code	
Country or region	

Service authorization and connectivity

Select the type of connection to contact your service provider. For a description of these methods that include security characteristics and configuration requirements, see “Choosing existing call-home servers to connect to service and support for this HMC” on page 70.

Table 18. Service authorization and connectivity

Characteristics	Entry field
Secure Sockets Layer (SSL) through the internet	_____
Virtual private network (VPN) through the internet	_____

Secure Sockets Layer (SSL) through the internet:

If you have an existing internet connection from your HMC, you can use it to call your service provider. You can connect directly to your service provider by using encrypted Secure Sockets Layer (SSL) by using the existing internet connection. Select **Use SSL Proxy** if you want to configure the use of encrypted SSL by using an indirect connection that uses an SSL Proxy.

Table 19. SSL

Characteristics	**Entry field**
Use SSL proxy? (yes/no)	
If yes, list information below:	
Address:	
Port:	
Authenticate with the SSL Proxy?	
If yes, list information below:	
User:	
Password:	

Internet connection Protocol used

For more information about the different internet Protocols, see "Configuring the HMC network types" on page 34.

___ IPv4

___ IPv6

___ IPv4 and IPv6

Virtual Private Network (VPN)

If you have an existing internet connection from your HMC, you can use it to call your service provider. You can connect directly to your service provider by virtual private network (VPN) by using the existing internet connection.

Note: If you select Virtual Private Network (VPN) through the internet, you cannot select any other options.

Call-home servers

Determine which HMCs you want to configure to connect to service and support as call-home servers. For more information about using multiple call-home servers, see "Using multiple call-home servers" on page 23.

___ This HMC

___ Another HMC

If you checked **Another HMC**, list the other HMCs that are configured as call-home servers here:

Table 20. Other HMCs that are configured as call-home servers

List of HMC host names or IP addresses that are configured as call-home servers

Extra Support Benefits

My Systems and Premium Search

Table 21. My Systems and Premium Search

Characteristics	Entry field
List your IBM ID	__________________________
List any additional IBM IDs	__________________________

To access valuable, customized support information in the My Systems and Premium Search sections of the Electronic Services website, Customers must register their IBM ID with this system. If you do not already have one, you can register for an IBM ID at: www.ibm.com/account/profile.

Note: IBM provides personalized web functions that use information that is collected by the IBM Electronic Service Agent application. To use these functions, you must first register on the IBM Registration website at http://www.ibm.com/account/profile.

To authorize users to use the Electronic Service Agent information to personalize the web functions, enter your IBM ID that you registered on the IBM Registration website. Go to http://www.ibm.com/support/electronic to see the valuable support information available to customers that register an IBM ID with their systems.

Configuring the HMC

Learn how to configure network connections, security, service applications, and some user preferences.

Depending on the level of customization you intend to apply to your HMC configuration, you have several options for setting up your HMC to suit your needs. The Guided Setup wizard is a tool on the HMC designed to ease the setup of the HMC. You can choose a fast path through the wizard to quickly create the recommended HMC environment, or you can choose to fully explore the available settings that the wizard guides you through. You can also perform the configuration steps without the aid of the wizard by Configuring the HMC by using the HMC menus.

Before you start, gather the required configuration information that you need to complete the steps successfully. See “Preparing for HMC configuration” on page 23 for a list of the required information. When you are finished preparing, ensure that you complete the “Preinstallation configuration worksheet for the HMC” on page 24 and then return to this section.

Configuring the HMC by using the menus

This section provides a complete list of all HMC configuration tasks, guiding you through the process of configuring your HMC. Choose this option if you prefer not to use the Guided Setup wizard.

You must restart your HMC for the configuration settings to take effect, so you might want to print this checklist and keep it with you as you configure your HMC.

This information contains references to tasks that are not included in this document. You can access the information centerIBM Power systems hardware information on the HMC or on the Web. On the HMC, IBM Knowledge Center can be accessed from the upper-right corner of the task bar. On the web, IBM Knowledge Center can be accessed at https://www.ibm.com/support/knowledgecenter.

This information contains references to tasks that are not included in this PDF. You can access additional support materials by referring to the **Additional Resources** section on the HMC Welcome page.

Prerequisites

Before you begin configuring the HMC using the HMC menus, be sure to complete the configuration preparation activity described in “Preparing for HMC configuration” on page 23.

Table 22. Manual HMC configuration tasks and where to find related information

Task	**Where to find related information**
1. Start the HMC.	"Starting the HMC" on page 33
2. Set the date and time.	
3. Change predefined passwords.	
4. Create additional users and return to this checklist when you have completed this step.	
5. Configure network connections.	"Configuring the HMC network types" on page 34
6. For HMC model 7063-CR2, you must configure the baseboard management controller (BMC) IP address.	"Configure BMC connectivity (7063-CR2)" on page 63
7. If you are using an open network and a fixed IP address, set identification information.	
8. If you are using an open network and a fixed IP address, configure a routing entry as the default gateway.	"Configuring a routing entry as the default gateway" on page 66
9. If you are using an open network and a fixed IP address, configure domain name services.	"Configuring domain name services" on page 66
10. If you are using a fixed IP address and have DNS enabled, configure domain suffixes.	"Configuring domain suffixes" on page 67
11. Configure your server to connect to IBM service and support and return to this checklist when you have completed this step.	"Configuring the local console to report errors to service and support" on page 68
12. Configure the Events Manager for Call Home.	"Configuring the Events Manager for Call Home" on page 72
13. Connect the managed system to a power source.	
14. Set passwords for the managed system, and each of the ASMI passwords (general and admin)	"Setting passwords for the managed system" on page 72
15. Access ASMI to set the date and time on the managed system.	
16. Start the managed system and return to this checklist when you have completed this step.	
17. Ensure that you have one logical partition on the managed system.	
18. Optional: add another managed system and return to this checklist when you have completed this step.	
19. Optional: If you are installing a new server with your HMC, configure the logical partitions and install the operating system.	
20. If you are not installing a new server at this time, perform optional postconfiguration tasks to further customize your configuration.	"Postconfiguration steps" on page 74

Starting the HMC

You can long in to the HMC and choose which language you want to be displayed in the interface. Use the default User ID `hscroot` and password `abc123` to log on to the HMC for the first time.

About this task

To start the HMC, do the following procedure:

Procedure

1. Turn on the HMC by pressing the power button.
2. If English is your language preference, continue with step 4.

   If your language preference is a language other than English, type the number **2** when you are prompted to change the locale.

   Note: This prompt times out in 30 seconds if you do not act.
3. Select the locale that you want to display from the list in the **Locale Selection** window, and click **OK**. The locale identifies the language that the HMC interface uses.
4. Click **Log on and launch the Hardware Management Console web application**.
5. Log in to the HMC with the following default user ID and password:

   ID: `hscroot`
   Password: `abc123`

   HMC Enhanced
   Displays the newer enhanced GUI with the enhanced PowerVM® features.

   HMC Classic
   Displays the standard GUI without the enhanced PowerVM features.

   Note: When the HMC is working as a DHCP server, the HMC uses the default password when it connects to the service processor for the first time.
6. Press Enter.

Changing the date and time

The battery-operated clock keeps the date and time for the Hardware Management Console (HMC). You might need to reset the console date and time if the battery is replaced, or if you physically move your system to a different time zone. Learn how to change the date and time for the HMC.

About this task

If you change the date and time information, the change does not affect the systems and logical partitions that the HMC manages.

To change the date and time for the HMC, complete the following steps:

Procedure

1. Ensure that you are a member of one of the following roles:
   - Super administrator
   - Service representative
   - Operator
   - Viewer
2. In the navigation area, click **Console Management**, and then select **Console Settings**.
3. In the content pane, click **Change Date and Time**.
4. If you select **UTC** in the **Clock** field, the time setting adjusts automatically for Daylight Saving Time in the time zone you select. Enter the date, time, and time zone, and click **OK**.

Results

Configuring the HMC network types

Configure your HMC so that it can communicate to the managed system, logical partitions, remote users, and service and support.

Configuring HMC settings to use an open network to connect to the managed system
Configure the HMC so that it can connect to and manage a managed system using an open network.

Before you begin

To configure the HMC network settings so that it can connect to the managed system using an open network, do the following:

Table 23. Configuring HMC settings to use an open network to connect to the managed system

Task	Where to find related information
1. Decide which interface you want to use for your managed system. **eth0** is preferred.	"Preinstallation configuration worksheet for the HMC" on page 24
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 36
3. Configure the Ethernet adapter by performing the following tasks:	
a. Set the media speed.	"Setting the media speed" on page 38
b. Select the open network type.	"Selecting a private or open network" on page 38
c. Set static addresses.	"Setting the IPv6 address" on page 64
d. Set the firewall.	"Changing HMC firewall settings" on page 64
e. Configure the default gateway.	"Configuring a routing entry as the default gateway" on page 66
f. Configure DNS.	"Configuring domain name services" on page 66
4. Configure additional adapters, if you have them.	
5. Test the connection between the managed server and the HMC.	"Testing the connection between the HMC and the managed system" on page 73

Configuring HMC settings to use a private network to connect to the managed system
Configure the HMC so that it can connect to and manage a managed system using a private network.

Before you begin

To configure the HMC network settings so that it can connect to the managed system using a private network, do the following:

Table 24. Configuring HMC settings to use a private network to connect to the managed system

Task	Where to find related information
1. Decide which interface you want to use for your managed system.	"Preinstallation configuration worksheet for the HMC" on page 24
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 36

Table 24. Configuring HMC settings to use a private network to connect to the managed system (continued)

Task	Where to find related information
3. Configure the HMC as a DHCP server.	"Configuring the HMC as a DHCP server" on page 38
4. Test the connection between the managed server and the HMC.	"Testing the connection between the HMC and the managed system" on page 73

Configuring HMC settings to use an open network to connect to logical partitions

Before you begin

To configure the HMC network settings so that it can connect to logical partitions using an open network, do the following:

Table 25. Configuring HMC settings to use an open network to connect to logical partitions

Task	Where to find related information
1. Decide which interface you want to use for your managed system.	"Preinstallation configuration worksheet for the HMC" on page 24
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 36
3. Configure the Ethernet adapter by performing the following tasks:	
a. Set the media speed.	"Setting the media speed" on page 38
b. Select the open network type.	"Selecting a private or open network" on page 38
c. Set static addresses.	"Setting the IPv6 address" on page 64
d. Set the firewall.	"Changing HMC firewall settings" on page 64
e. Configure the default gateway.	"Configuring a routing entry as the default gateway" on page 66
f. Configure DNS.	"Configuring domain name services" on page 66
4. Configure additional adapters, if you have them.	
5. Test the connection between the managed server and the HMC.	"Testing the connection between the HMC and the managed system" on page 73

Configuring HMC settings to use an open network to connect to remote users

Before you begin

To configure the HMC network settings so that it can connect to remote users using an open network, do the following:

Table 26. Configuring HMC settings to use an open network to connect to remote users

Task	Where to find related information
1. Decide which interface you want to use for your managed system.	"Preinstallation configuration worksheet for the HMC" on page 24
2. Identify the Ethernet ports for your HMC.	"Identifying the Ethernet port that is defined as eth0" on page 36

Table 26. Configuring HMC settings to use an open network to connect to remote users (continued)

Task	Where to find related information
3. Configure the Ethernet adapter by performing the following tasks:	
a. Set the media speed.	“Setting the media speed” on page 38
b. Select the open network type.	“Selecting a private or open network” on page 38
c. Set static addresses.	“Setting the IPv6 address” on page 64
d. Set the firewall.	“Changing HMC firewall settings” on page 64
e. Configure the default gateway.	“Configuring a routing entry as the default gateway” on page 66
f. Configure DNS.	“Configuring domain name services” on page 66
g. Configure suffixes.	“Configuring domain suffixes” on page 67
4. Configure additional adapters, if you have them.	

Configuring HMC call-home server settings

Before you begin

To configure the HMC call-home server settings so that problems can be reported, do the following:

Table 27. Configuring HMC call-home server settings

Task	Where to find related information
1. Be sure you have all the required customer information	“Preinstallation configuration worksheet for the HMC” on page 24
2. Configure this HMC to report errors or choose an existing call-home server to report errors	“Configuring the local console to report errors to service and support” on page 68 “Choosing existing call-home servers to connect to service and support for this HMC” on page 70
3. Verify that your call-home configuration is working	“Verifying that your connection to service and support is working” on page 70
4. Authorize users to view collected system data	“Authorizing users to view collected system data” on page 71
5. Schedule transmission of system data	“Transmitting service information” on page 71

Identifying the Ethernet port that is defined as eth0
Your Ethernet connection to the managed server must be made by using the Ethernet port that is defined as `eth0` on your HMC.

If you did not install any additional Ethernet adapters in the PCI slots on your HMC, then the primary-integrated Ethernet port is always defined as `eth0` or `eth1` on your HMC, if you intend to use the HMC as a DHCP server for your managed systems.

If you install extra Ethernet adapters in the PCI slots, then the port that is defined as `eth0` depends on the location and type of Ethernet adapters that are installed.

Note: The following general rules might not apply for all configurations.

The following table describes the rules for Ethernet placement by HMC type.

Table 28. HMC types and associated rules for Ethernet placement

HMC type	Rules for Ethernet placement
Rack-mounted HMCs with two integrated Ethernet ports.	The HMC supports only one extra Ethernet adapter. • If an extra Ethernet adapter is installed, then that port is defined as `eth0`. In this case, the primary-integrated Ethernet port is then defined as `eth1`, and the secondary integrated Ethernet port is defined as `eth2`. • If the Ethernet adapter is a dual port Ethernet adapter, then the port that is labeled Act/Link A is `eth0`. The port that is labeled `Act/link` B is `eth1`. In this case, the primary-integrated Ethernet port is then defined as `eth2`, and the secondary integrated Ethernet port is defined as `eth3`. • If no adapters are installed, then the primary-integrated Ethernet port is defined as `eth0`.
Stand-alone models with a single integrated Ethernet port.	The definitions depend upon the type of Ethernet adapter that is installed: • If only one Ethernet adapter is installed, then that adapter is defined as `eth0`. • If the Ethernet adapter is a dual port Ethernet adapter, then the port that is labeled `Act/link` A is `eth0`. The port that is labeled `Act/link` B would be `eth1`. In this case, the primary-integrated Ethernet port is then defined as `eth2`. • If no adapters are installed, then the integrated Ethernet port is defined as `eth0`. • If multiple Ethernet adapters are installed, see "Determining the interface name for an Ethernet adapter" on page 37.

Determining the interface name for an Ethernet adapter
If you configure the HMC as a DHCP server, that server can operate only on the network interface card (NIC) connectors that the HMC identifies as `eth0` and `eth1`. You might also need to determine which NIC connector you need to plug the Ethernet cable into. Learn more about determining which NIC connectors the HMC identifies as `eth0` and `eth1`.

About this task

To determine the name the HMC has assigned to an Ethernet adapter, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change Network Settings**.
3. From the **Change Network Settings** window, click the **LAN adapters** tab. The following example entry shows that this Ethernet port is identified as `eth0`: `Ethernet eth0 52:54:00:fa:b6:8e (<IP address of HMC>)`.
4. Record your results. If you need to view or change the LAN adapter settings, click **Details**.
5. Click **OK**.

Setting the media speed
Learn how to specify the media speed that includes the speed and duplex mode of the Ethernet adapter.

Before you begin

The default for the HMC adapter settings is **Autodetection**. If this adapter is connected to a LAN switch, you must match the switch port settings. To set the media speed and duplex, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. In the local area network (LAN) information section, select **Autodetection** or the appropriate media speed and duplex combination.
6. Click **OK**.

Selecting a private or open network
A *private service network* consists of the Hardware Management Console (HMC) and the managed systems. A private service network is restricted to consoles and the systems they manage, and is separate from your company network. An *open network* consists of your private service network and your company network. An open network might contain network endpoints in addition to consoles and managed systems, and might span across multiple subnets and network devices.

About this task

To select a private or public network, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **LAN Adapter** tab.
6. In the local area network information page, select **Private** or **Open**.
7. Click **OK**.

Configuring the HMC as a DHCP server
Dynamic Host Configuration Protocol (DHCP) provides an automated method for dynamic client configuration.

To configure the Hardware Management Console (HMC) as a DHCP server, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Customize Network Settings window opens.
3. Select the LAN adapter that you want to work with and click **Details**.
4. Select **Private** and then select the network type.
5. In the DHCP Server section, select **Enable DHCP Server** to enable the HMC as a DHCP server.

   Note: You can configure the HMC to be a DHCP server only on a private network. If you use an open network, the option to select the **Enable DHCP** is not available.
6. Enter the address range of the DHCP server.

7. Click **OK**.

If you configured your HMC to be a DHCP server on a private network, you must verify that your HMC DHCP private network is configured correctly. For information about connecting your HMC to a private network, see "Selecting a private or open network" on page 38.

For more information, see " HMC as a DHCP server" on page 17.

Managing the system by using OpenBMC-based HMC (7063-CR2)
IBM Power Systems servers use a baseboard management controller (BMC) for system service management, monitoring, maintenance, and control. The BMC also provides access to the system event log files (SEL). The BMC is a specialized service processor that monitors the physical state of the system by using sensors. A system administrator or service representative can communicate with the BMC through an independent connection. The OpenBMC tool provides a communication method to the BMC, by using a command-line interface. The OpenBMC tool can be used either from a remote Linux system, or from the host operating system console window. The OpenBMC tool can be connected remotely to the BMC by using a configured Ethernet port. You can connect your server to a monitor by using the VGA port at the rear of the server.

Managing the system by using the OpenBMC tool
Learn how to configure and manage your system by using the OpenBMC tool.

Downloading and installing the OpenBMC tool
Learn how to download and install the OpenBMC tool.

About this task

To download and install the OpenBMC tool, complete the following steps:

Procedure

1. Go to the IBM Support Portal.
2. In the search field, type: `Scale-out LC System Event Log Collection Tool`.
3. Click the **Scale-out LC System Event Log Collection Tool** entry and follow the instructions to install and run the OpenBMC tool.

Basic commands and functionality of the OpenBMC tool
The OpenBMC tool provides support for working with system event logs, updating system firmware, identifying the system, powering off the system, and other service-related functions.

OpenBMC tool top-level options
Learn more about the top-level options for the OpenBMC tool commands.

About this task

- `-H`: Host name or IP address of the BMC.
- `-U`: User name to log in with.
- `-A`: Provides a prompt to ask for the password.
- `-P`: Password for the user name.
- `-j`: Change output format to JSON.
- `-t`: Location of the policy table to use.
- `-T`: Provides time statistics for logging in, running the command, and logging out.
- `-V`: Displays current version of the OpenBMC tool.

System event log commands
Learn more about system event log commands for the OpenBMC tool.

Procedure

- To print a list of the system event logs in a readable format, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel print
  ```

- To list the system event logs in raw data, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel list
  ```

- To change the status of a system event log to resolved, use the following command:

  `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel resolve -n x`, where *x* is the system event log number.

- To collect all service data including system event logs, use the following command:

  `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> collect_service_data`.

- To clear gard records for disable hardware, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> gardclear
  ```

- To clear the alert logs of entries, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel clear
  ```

System firmware update command
Learn more about the system firmware update command.

Procedure

- To update the system firmware, use the following command:

  `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> firmware flash <bmc or pnor> -f xxx.tar`, where *bmc* or *pnor* is the type of image you wish to flash to the system.

  Note: If you are not in the same folder as the TAR file, you must include the full path to the folder where the file resides.

- To activate a firmware image that is available in the BMC, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> firmware activate <firmware image ID>
  ```

System identify commands
Learn more about the system identify commands.

Procedure

- To activate the blue system identify LED, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis identify on
  ```

- To turn off the blue system identify LED, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis identify off
  ```

- To check the status of the blue system identify LED, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
chassis identify status
```

System power on and power off commands
Learn more about the system power on and power off commands.

Procedure

- To check the power status of the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  chassis power status
  ```
- To power on the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  chassis power on
  ```
- To power off the system normally, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  chassis power softoff
  ```
- To power off the system immediately, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  chassis power hardoff
  ```

System sensor commands
Learn more about the system sensor commands.

Procedure

- To display a list of all monitoring sensors, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  sensors print
  ```

  or

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  sensors list
  ```

System FRU commands
Learn more about the system FRU commands.

Procedure

- To display a list of all inventory items, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  fru print
  ```

  or

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  fru list
  ```
- To display the known status of all FRU items, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  fru status
  ```

  Note: The FRU item must be designated as a replaceable FRU by the BMC.
- To automate the review of FRU status commands and to determine if there is a performance impact on the system, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
health_check
```

Note: This command does not guarantee a healthy system as there can be system event logs entries that are not associated with the inventory items.

System BMC reset commands
Learn more about the system BMC reset commands.

Procedure

- To do a warm reset of the BMC remotely and without an AC cycle, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  bmc reset warm
  ```
- To do a cold reset of the BMC remotely and without an AC cycle, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  bmc reset cold
  ```

System dump commands
Learn more about the system dump commands.

Procedure

- To create a new dump file, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump create
  ```
- To list all dump files in the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump list
  ```
- To delete a specific dump file from the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump delete -n <dump file entry>
  ```
- To delete all dump files from the system, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump delete all
  ```
- To retrieve a specific dump file, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump retrieve -n <dump file entry>
  ```
- To retrieve a dump file and save it to specific directory, use the following command:

  ```
  openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name>
  dump retrieve -s <location to save dump file>
  ```

  Note: If you do not specify a location, the file is saved in the OS where the command is run in the temp directory.

Enabling and disabling local BMC user accounts
Learn more about the **local_users**commands.

About this task

The local user accounts on the BMC, such as root, can be disabled, queried, and re-enabled with the **local_users** sub-command.

Note: After disabling local users, the LDAP user needs to be available for further interaction with the BMC, including enabling local users by using OpenBMC tool.

Procedure

- To view current local user account status, use the following command:

  ```
  openbmctool <connection options> local_users queryenabled
  ```
- To disable all local user accounts, use the following command:

  ```
  openbmctool <connection options> local_users disableall
  ```
- To enable all local user accounts, use the following command:

  ```
  openbmctool <connection options> local_users enableall
  ```

Remote logging by using rsyslog
Learn more about the remote logging commands.

About this task

The BMC can stream out local logs (that go to the systemd journal) by using RSYSLOG. The BMC sends everything in the logs. Any kind of filtering and appropriate storage has to be managed on the rsyslog server.

Procedure

- To configure the rsyslog server for remote logging, use the following command:

  ```
  openbmctool <connection options> logging remote_logging_config -a <IP
  address> -p <port>
  ```

  Note: The IP address and port are for the remote rsyslog server. After the command is run, the remote rsyslog server starts to receive logs from the BMC.
- To disable remote logging, use the following command:

  ```
  openbmctool <connection options> logging remote_logging disable
  ```

  Note: Disable remote logging before you switch remote logging from an existing remote server to a new one.
- To view the remote logging configuration, use the following command:

  ```
  openbmctool <connection options> logging remote_logging view
  ```

  Note: This command prints out the IP address and port of the remote rsyslog server in JavaScript Object Notation (JSON) format.
- To turn REST API logging on, use the following command:

  ```
  openbmctool <connection options> logging rest_api on
  ```
- To turn REST API logging off, use the following command:

  ```
  openbmctool <connection options> logging rest_api off
  ```

  Note: REST API logging is turned off by default.

Certificate management
Learn more about the certificate management commands.

About this task

You can replace the existing certificate and private key file with another (possibly CA signed) certificate and private key file. You can install server, client, and root certificates.

Procedure

- To update the HTTPS server certificate, use the following command:

  ```
  openbmctool <connection options> certificate update server https -f <File>
  ```

  Note: The <File> is the privacy-enhanced mail (PEM) file that contains both the certificate and the private key.
- To update the LDAP client certificate, use the following command:

  ```
  openbmctool <connection options> certificate update client ldap -f <File>
  ```

  Note: The <File> is the PEM file that contains both the certificate and the private key.
- To update the LDAP root certificate, use the following command:

  ```
  openbmctool <connection options> certificate update authority ldap -f <File>
  ```

  Note: The <File> is the PEM file that contains only the certificate.
- To delete the HTTPS server certificate, use the following command:

  ```
  openbmctool <connection options> certificate delete server https
  ```

  Note: Deleting a certificate creates and installs a new self-signed certificate.
- To delete the LDAP client certificate, use the following command:

  ```
  openbmctool <connection options> certificate delete client ldap
  ```
- To delete the LDAP root certificate, use the following command:

  ```
  openbmctool <connection options> certificate delete authority ldap
  ```

  Note: Deleting the root certificate can cause an LDAP service outage.

LDAP configuration
Learn more about the LDAP configuration commands.

About this task

In the BMC, LDAP is used for remote authentication. The BMC does not support remote user-management functionality. The BMC supports both secure and non-secure LDAP configuration.

Procedure

- To create the LDAP configuration (non-secure), use the following command:

  ```
  openbmctool.py <connection options> ldap enable --uri="ldap://
  <ldap server IP/hostname>" --bindDN=<bindDN> --baseDN=<basDN> --
  bindPassword=<bindPassword> --scope="sub/one/base" --serverType="OpenLDAP/
  ActiveDirectory"
  ```

  Note: Configuring a fully qualified domain name or hostname in the `uri` parameter requires the domain name system (DNS) server to be configured on the BMC.
- To create the LDAP configuration (secure), use the following command:

  ```
  openbmctool.py <connection options> ldap enable --uri="ldaps://
  <ldap server IP/hostname>" --bindDN=<bindDN> --baseDN=<basDN> --
  bindPassword=<bindPassword> --scope="sub/one/base" --serverType="OpenLDAP/
  ActiveDirectory"
  ```

  Notes:

  1. It is common to encounter the following error when you run the above `openbmctool.py` command string:

     xyz.openbmc_project.Common.Error.NoCACertificate

This error means that the BMC client needs to verify that the LDAP server certificate is signed by a known certification authority (CA). An administrator needs to upload the CA certificate to the BMC to resolve this error.

2. The OpenBMC tool does not support individual LDAP configuration property updates. To update a single property, the administrator must recreate the LDAP configuration with the changed values.

- To delete the LDAP configuration, use the following command:

  ```
  openbmctool.py <connection options> ldap disable
  ```

  Note: The root user must be enabled before you run the command, otherwise the BMC is not accessible. To enable all local user accounts, see Enabling and disabling local user accounts.

- To add privilege mapping use the following command:

  ```
  openbmctool.py <connection options> ldap privilege-mapper create --
  groupName=<groupName> --privilege="priv-admin/priv-user"
  ```

- To delete privilege mapping, use the following command:

  ```
  openbmctool.py <connection options> ldap privilege-mapper delete --
  groupName=<groupName>
  ```

- To list privilege mapping, use the following command:

  ```
  openbmctool.py <connection options> ldap privilege-mapper list
  ```

  The normal workflow for LDAP configuration is in the following order:

  1. Configure the DNS server.
  2. Configure LDAP.
     a. Configure the CA certificate for secure LDAP configuration.
     b. Create LDAP configuration with local user.
  3. Configure user privilege.

  Notes:

  1. If you login with LDAP credentials and have not added privilege mapping for the LDAP credentials, then you will get the following error message:

     403, 'LDAP group privilege mapping does not exist'.

     You can avoid this error by adding privilege mapping.

  2. The following error message might mean that user lacks sufficient privileges on the BMC:

     Insufficient privileges

     You can avoid this error by adding privilege mapping.

  3. After you setup the LDAP, the OpenBMC tool connection options work with both LDAP and local users.

Network configuration
Learn more about the network configuration commands.

Procedure

- To enable DHCP, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network enableDHCP -I
  <Interface name>
  ```

- To disable DHCP, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network disableDHCP -I
  <Interface name>
  ```

- To get the host name, use the following command:

```
openbmctool.py -H <BMC_IP> -U root -P <root password> network getHostName
```

- To set the host name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setHostName -H
  <host name>
  ```

- To get the domain name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getDomainName
  -I <Interface name>
  ```

- To set the domain name, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setDomainName
  -I <Interface name> -D DomainName1,DomainName2,..
  ```

- To get the media access control (MAC) address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getMACAddress
  -I <Interface name>
  ```

- To set the MAC address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setMACAddress
  -I <Interface name> -MA xx:xx:xx:xx:xx
  ```

- To get the default gateway, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getDefaultGW
  ```

- To set the default gateway, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setDefaultGW
  -GW <default gw>
  ```

- To view the current network configuration, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network view-config
  ```

- To get the network time protocol (NTP), use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getNTP -I
  <Interface name>
  ```

- To set the NTP, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setNTP -I
  <Interface name> -N NTP1,NTP2,...
  ```

- To get the domain name system (DNS), use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getDNS -I
  <Interface name>
  ```

- To set the DNS, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network setDNS -I
  <Interface name> -d DNS1,DNS2,...
  ```

- To get the IP address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network getIP -I
  <Interface name>
  ```

- To set the IP address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network addIP -a
  <ADDRESS> \-gw <GATEWAY> -l <PREFIXLENGTH> -p <protocol type> -I <Interface
  name>
  ```

- To delete the IP address, use the following command:

  ```
  openbmctool.py -H <BMC_IP> -U root -P <root password> network rmIP -I
  <Interface name> -a <ADDRESS>
  ```

- To enable a virtual local area network (VLAN), use the following command:

  ```
  openbmctool.py <connection options> network addVLAN -I <Interface name> -n
  <IDENTIFIER>
  ```

- To disable a virtual local area network (VLAN), use the following command:

  ```
  openbmctool.py <connection options> network deleteVLAN -I <Interface name>
  ```

- To view the DHCP configuration properties, use the following command:

  ```
  openbmctool.py <connection options> network viewDHCPConfig
  ```

- To configure the DHCP properties, use the following command:

  ```
  openbmctool.py <connection options> network configureDHCP -d <DNSENABLED> -n
  <HOSTNAMEENABLED> -t <NTPENABLED> -s <SENDHOSTNAMEENABLED>
  ```

  Note: `DNSENABLED`, `HOSTNAMEENABLED`, `NTPENABLED`, and `SENDHOSTNAMEENABLED` are boolean values (true or false).

- To reset the network settings to the factory defaults, use the following command:

  ```
  openbmctool.py <connection options> network nwReset
  ```

  Note: Reset settings are applied after the rebooting of the BMC.

Managing the system by using the IPMI
Learn how to configure and manage your system by using the Intelligent Platform Management Interface (IPMI).

Common IPMI commands
You can use **`IPMI`** commands to perform various managing tasks for your system.

Table 29. Common IPMI commands

Command option	Description
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `chassis power on`	Powers on the server.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `chassis power off`	Powers off the server.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `chassis status`	Checks the server status.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `chassis power cycle`	Power cycle the server.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `sol activate`	Activates SOL system console.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `sol deactivate`	Deactivates SOL system console.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `sel list`	Returns an error log.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `sdr list`	Lists status of all sensors.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `sol set retry-interval` *`value`*	Sets the default retry-interval value in milliseconds.
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `fru print`	Prints the FRU information.

Table 29. Common IPMI commands (continued)

Command option	Description
`ipmitool -I lanplus -H` *`myserver.example.com`* `-P` *`mypass`* `user list`	Lists the IPMI users.

Configuring the BMC IP address
Dynamic Host Configuration Protocol (DHCP) is the default network setup for the BMC in the 7063-CR2 HMC. To enable your network connection, you can connect to your system and use the Petitboot bootloader interface to configure the IP address of the BMC. If you do not plan to use DHCP, you can also set up a static IP address. Alternatively, you can use the HMC GUI to enable your network connection by navigating to **HMC Management** > **Console Settings** > **Change BMC/IPMI Network Settings**.

Before you begin

You must connect the network cable and a VGA monitor before you access the Petitboot bootloader interface.

If you encounter any problems in accessing the Petitboot bootloader interface, see Resolving a BMC access problem.

About this task

To use the Petitboot bootloader interface to set up or enable the network interface of the BMC, complete the following steps:

Procedure

1. Power on the server by pressing the power button on the front of the system. The system powers on to the Petitboot bootloader menu.

   Note: The boot process takes about 1 to 2 minutes to complete.

   When Petitboot loads, the monitor activates. Press any key to interrupt the boot process.
2. At the Petitboot bootloader main menu, select **Exit to Shell**.
3. Run the following command: `ipmitool lan print 2`. If this command returns an IP address, verify that is correct. To set a static IP address, follow these steps:

   Notes:

   - The following two LAN interfaces are available to BMC:
     - Shared interface is LAN1
     - Dedicated interface is LAN2
   - The `ipmitool lan print 2` command cannot display more than one IP address. To avoid this situation, you can set the static IP address to a different subnet from the default zero configuration networking IP address.

   a. Set the mode to static by running the following command: `ipmitool lan set 2 ipsrc static`.

   b. Set your IP address by running the following command: `ipmitool lan set 2 ipaddr ip_address`, where *ip_address* is the static IP address that you want to assign to this system.

   c. Set your netmask by running the following command: `ipmitool lan set 2 netmask netmask_address`, where *netmask_address* is the netmask for the system.

   d. Set your gateway server by running the following command: `ipmitool lan set 2 defgw ipaddr gateway_server`, where *gateway_server* is the gateway for this system.

   e. Confirm the IP address by running the following command: `ipmitool lan print 2`.

Performing a factory reset
Learn how to perform a factory reset on the system.

The factory reset function can take up to 15 minutes to complete. When the LED on the power button starts flashing, the system is ready to start again. Perform the factory reset with the host powered off. If you perform the factory reset while the host is running, the system shuts down immediately and restarts the BMC. If the BMC is on a static network, you must manually power on the system with the physical power button.

To perform a factory reset, run the following command:

```
ipmitool -I lanplus -U <username> -P <password> -H <BMC_IP or Hostname> raw 0x3A 0x11
```

Note: The system does not send a validation response. The following system output is normal:

```
Unable to send RAW command (channel=0x0 netfn=0x3a lun=0x0 cmd=0x11)
```

If you forgot the password of the BMC, you can run the following command while the host is running to perform a factory reset and to restore the default password:

```
ipmitool raw 0x3A 0x11
```

You must set up and configure the BMC IP address after performing the factory reset. For more information, see "Configuring the BMC IP address" on page 48.

Risks of using IPMI on IBM Power systems and OpenPower Systems
Various risks that are associated with the Intelligent Platform Management Interface (IPMI) have been identified and documented in the information technology (IT) security community.

IBM Power systems and OpenPower Systems provide IPMI access by default. A subset of these identified risks is applicable to IBM servers.

Vulnerability Details

The IPMI service can become unresponsive after it receives and rejects multiple authentication attempts. You might receive a `insufficient resources for session` message if you use the IPMI immediately after the failed authentication attempts. This situation lasts for a few seconds and normal service is restored afterward.

Important: Repeated authentication failures can cause denial of service.

A list of common vulnerabilities and exposures (CVE) is listed in Table 30 on page 49.

Table 30. Common vulnerabilities and exposures

CVE ID	Description
CVE-2013-4037	The Remote Authenticated Key-Exchange Protocol (RAKP), which is specified by the IPMI standard for authentication, has flaws. Although the system does not allow the use of null passwords, a hacker might reverse engineer the RAKP transactions to determine a password. The authentication process for IPMI requires the management controller to send a hash of the requested password of the user to the client before the client authenticates. This process is a key part of the IPMI specification. The password hash can be broken by using an offline brute force or dictionary attack.

Table 30. Common vulnerabilities and exposures (continued)

CVE ID	Description
CVE-2013-4031	IBM Power systems and OpenPower Systems are preconfigured with one IPMI user account, which has the same default login name and password on all affected systems. If a malicious user gains access to the IPMI interface by using this preconfigured account, the user can power off or on, or restart the host server, and create or change user accounts possibly preventing legitimate users from accessing the system. On OpenPower Systems, the default IPMI user name is `root`. Additionally, if a user fails to change the default user name and password on each of the systems that is deployed, the user has the same login information for each of those systems.
CVE-2013-4786	The IPMI 2.0 specification supports RMCP+ Authenticated Key-Exchange Protocol (RAKP) authentication, which allows remote attackers to obtain password hashes and conduct offline password guessing attacks by obtaining the hash-based message authentication code (HMAC) from a RAKP message 2 response from a BMC.

Configuration options and best practices

- Change the preconfigured user name and password when the server is deployed. This action prevents unauthorized users from gaining access to the system through the preconfigured user account.
- Do not disable the IPMI access for the user whose access credentials have been shared with the HMC via the **Console Inband Communication Credentials** task.

  Note: To launch the **Console Inband Communications Credentials** task, complete the following steps:

  1. In the navigation area, click **Console Management**, and then select **Console Settings**.
  2. In the content pane, click **Console Inband Communications Credentials**.
  3. From the **Console Inband Communications Credentials** window, you can set the inband BMC credentials or modify an expired password for previously set inband BMC credentials for the HMC.
- If a user is not managing a server by using the IPMI, you can configure the system to disallow IPMI network access from the user accounts. This task can be accomplished by using the IPMItool utility or a similar utility for managing and configuring the IPMI management controllers. You can use the following IPMItool command to disable the network access for an IPMI user:

  ```
  ipmitool channel setaccess 1 #user_slot# privilege=15
  ```

  Note: Replace #user_slot# in the command with the actual slot number (1 - 12) and repeat for each configured user.

  This example shows the command when it is run directly on the server. If the IPMItool command is run remotely over the network, or if a different utility is used, the command might be different. See the documentation for the utility that you are using to determine the correct command syntax. Disallowing IPMI network access removes the ability to use the weakness that is present in the IPMI RAKP protocol to discover user account credentials.
- Use strong passwords that are at least 16 characters long with a mixture of upper and lowercase letters, numbers, and special characters. By using more complex passwords, it makes it more difficult for malicious users to discover valid user credentials.
- Keep the management network separate from the public network. Keeping the management network separate lessens security exposures by reducing the number of individuals who can access the systems.

Managing the system by using the OpenBMC GUI
Learn how to manage and configure your system by using the OpenBMC GUI.

Logging on to the OpenBMC GUI
Learn how to log on to the OpenBMC GUI.

To log on to the OpenBMC GUI, complete the following steps:

1. Open a supported web browser. In the address bar, enter the IP address of the BMC that you want to connect to. For example, you can use the format `https://<BMC IP>` in the address bar of the web browser.
2. From the **OpenBMC logon** window, enter the **Host** address of the BMC and the **Username** and **Password** that is assigned to you.

   Note: The default user ID is `root` and the default password is `0penBmc`.

   If you are using firmware level OP940.01, or later, the root password is expired by default. You must change the default password before you can access the BMC. For more information about changing the expired default password, see "Setting the password" on page 51.

   If you forgot your password, you can perform a factory reset of the system to restore the default password. To reset the system, see "Performing a factory reset" on page 49.
3. Click **Log in**.

Setting the password
Learn how to change and set the password for your **root** account and to help secure the system.

Improved BMC password policy

The baseboard management controller (BMC) **root** password must be set on first use for newly manufactured systems or after performing a factory reset of the system. This policy change helps to enforce that the BMC is not left in a state with a well-known password.

In firmware level OP940.01, and later, the root password is expired and must be changed before you can access the functions of the BMC. However, if you are upgrading the firmware level from a previous OpenBMC firmware level or if you are performing an operational installation, you do not have to change the password.

The default user ID is `root` and the default password is `0penBmc`. You can use the web application, the Redfish REST APIs, the OpenBMC tool command to change the password. You can also use the **Console Inband Communications Credentials** task in the HMC GUI to change the expired password.

Note: To launch the **Console Inband Communications Credentials** task, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Console Inband Communications Credentials**.
3. From the **Console Inband Communications Credentials** window, you can set the inband BMC credentials or modify an expired password for previously set inband BMC credentials for the HMC.

After changing the password, you can access the BMC with your usual interface. To change the password, you must first access the account with the correct credentials, and then use the password change function. If you attempt to access the BMC with an expired password, you must change the password before accessing other functions.

- To change your expired password by using the web interface, enter `https://<BMC_IP>` into a web browser and then enter the access credentials of the BMC. The web interface prompts you to enter a new password.
- To change your expired password through a network interface, you can use Redfish APIs. For instructions, see "Managing the system by using DMTF Redfish APIs" on page 59.

- To change your expired password by using the OpenBMC tool, run the `openbmctool set_password` subcommand. For example,

  ```
  openbmctool.py -H <BMC IP address or BMC host name> -U <username> -P <password> set_password
  -p <new password>
  Attempting login...
  200
  User root has been logged out
  ```

  Where `200` is the response status that indicates success.

Note: The system might take up to 5 minutes to update the new password on the BMC. If you have trouble accessing your account, wait for 5 minutes and try again.

Also, with firmware level OP940.01, the BMC factory reset function resets the BMC password back to its default value and causes the default password to expire. This function means that after you perform the factory reset, you must change the password before you can access the BMC (even if you upgraded from an older firmware level).

To increase account security of the system, the administrator must complete the following steps:

1. Set a strong password for the root account. Strong passwords have at least 15 characters and include non-alphabetic characters. Initially, the password must not exceed 20 characters. Passwords can be changed later to a length greater than 20 characters, but IPMI access will be removed. Avoid using the **root** account, as the **root** account has more access to the BMC than an **Administrator** account. The root account can present a security risk if it is used incorrectly or maliciously. Use the root account only when it is required.
2. Create a separate account for each entity to manage the system. For example, you can create an **Administrator** account for yourself and for xCat, and create an **Operator** account for your staff. You can use the web interface or Redfish APIs to create a new account. When you create a new account, carefully consider which privilege role to assign to the user. Always use the least privilege role that is required.

   - To create a new account by using the web interface, see "Local users" on page 57.
   - To create a new account by using the Redfish APIs, see "Managing the system by using DMTF Redfish APIs" on page 59.

   If your BMC is using Lightweight Directory Access Protocol (LDAP), you can add users to the LDAP server.
3. Log off from the root account and switch to your personal **Administrator** account.

To increase the security of the system, the administrator can optionally configure access to the LDAP server. For more information, see "Basic commands and functionality of the OpenBMC tool " on page 39.

Dashboard

The dashboard displays the overall information about the server and the BMC.

The following options are available on the title bar (located in the top portion of the dashboard):

- **Server information**: Displays the server name and BMC IP address.
- **Server health**: Displays the status of the server.
- **Server power**: Displays whether the server is powered on, powered off, or in an error state.
- **Date last refreshed**: Displays the date and time that the information was last refreshed. The time zone of the user is determined by the web browser.
- **Refresh**: Click **Refresh** to refresh the information.

The following menus are available from the menu pod (located in the left portion of the dashboard):

- **Server overview**
- **Server health**
- **Server control**
- **Server configuration**

- **Users**

Server overview
Learn about the options that are available from the **Server overview** task.

From the **Server overview** window, you can choose from any of the following available options:

- **Server information**: Displays the model, manufacturer, firmware version, and serial number of the server.
- **BMC information**: Displays the host name, BMC IP address, firmware version, and MAC address of the BMC.
- **Power information**: Displays the power consumption and power cap.
- **High priority events**: View any high priority events. Click **Refresh** to reload the information that is displayed here.
- **BMC time**: Displays the BMC time in the time zone of the user, which is determined by the web browser.
- **Turn on server LED**: Turn on or turn off the server LED.
- **Launch serial over LAN console**: Launches the Serial over LAN (SoL) console.
- **Edit network settings**: Edit the network settings.

Server health
Learn about the tasks that are available from the **Server health** menu.

From this menu, you can choose from any of the following available tasks:

Event log
View all events from the BMC.

You can view and filter event log files from the BMC. From the **Event log** window, you can perform the following actions:

- Search through event logs by entering keywords and clicking **Search**.
- Filter the event logs by severity (**All**, **High**, **Medium**, or **Low**). You can select multiple severity levels.
- Filter the event logs by date range.
- Filter the event logs by event status (**All events**, **Resolved events**, and **Unresolved events**).
- Click any of the events that are listed to expand the event log file for more information. You can click **Copy** to copy the information to the clipboard.
- Select multiple event logs by clicking the checkbox next to event log. After you select the event logs, you can delete the logs by clicking **Delete** and then clicking **Yes** in the confirmation message. You can also mark event logs as read by clicking **Mark as resolved**.

Hardware status
View the hardware status and associated events of all hardware in the server.

You can view the hardware status of various hardware components in your server. Click any of the hardware components to expand the view for more information. You can search for specific hardware components by using the **Filter Hardware Components** search feature and then clicking **Filter**. You can also export the data by clicking **Export**.

Sensors
View all sensors that are present in the system.

You can view and filter sensors from the BMC. From the **Sensors** window, you can perform the following actions:

- Search and filter for specific sensors by using the **Search** feature and then clicking **Filter**.
- Filter sensors by severity (**All**, **Critical**, **Warning**, or **Normal**).
- Export the sensor data by clicking **Export**.

Server control
Learn about the tasks that are available from the **Server control** menu.

From this menu, you can choose from any of the following available tasks:

Server power operations
Learn how to view current server status and select power operations.

To update the **Host OS boot settings**, complete the following steps:

Note: It is not recommended to modify the **Host OS boot settings** for 7063-CR2 HMC, unless instructed by IBM Support.

1. Select the boot setting override type from the **Boot setting override** menu.
2. You can optionally select **Enable one time boot**.
3. Select whether to enable or disable the TPM policy by selecting **On** or **Off**.
4. Click **Save**.

To restart the server, complete the following steps:

1. Select the type of reboot from **Operations** > **Reboot server**:
   - Orderly: Operating system shuts down first and then the server reboots.
   - Immediate: Server reboots without the operating system shutting down. This might cause data corruption.
2. Click **Reboot**.

To shutdown the server, complete the following steps:

1. Select the type of shutdown from **Operations** > **Shutdown server**:
   - Orderly: Operating system shuts down first and then the server reboots.
   - Immediate: Server reboots without the operating system shutting down. This might cause data corruption.
2. Click **Shutdown**.

Manage power usage
Learn how to view the power consumption of the server and set a power cap.

Note: It is recommended to set the **power cap** to **Off** for 7063-CR2 HMC. By default, the **power cap** is set to **Off**.

To set a power cap, complete the following steps:

1. From the **Server power cap setting** section, set the **power cap** to **On**.
2. Specify the number of watts to keep the server power consumption at or below the specified value.
3. Click **Save settings**.

You can turn off the power cap by setting the **power cap** to **Off** and clicking **Save settings**.

Server LED
Learn how to turn on and turn off the server light-emitting diode (LED).

You can turn on or turn off the server LED by clicking the toggle switch to either **On** or **Off**.

Note: If the server has a liquid crystal display (LCD), you can use this control to display text (**On**) or not to display text (**Off**) on the LCD.

Reboot BMC
Learn how to restart the BMC and view the current BMC boot status.

Click **Reboot BMC** to restart the BMC.

Note: When you restart the BMC, your web browser looses connection with the BMC for several minutes. When the BMC is back online, you must log in again. If the **Log in** button is not available after you restart the BMC, close your web browser. Then, reopen the web browser and enter your BMC IP address.

Serial over LAN console
Learn how to view information over the serial port of the server.

You can launch the Serial over LAN (SoL) console that displays the output of the serial port of the server.

KVM
Learn how to launch the remote keyboard, video, and mouse (KVM) console.

You can launch the KVM console from this task and interact with the remote system.

Virtual media
Learn how to start a session by using a virtual media device.

To start a session, complete the following steps:

1. Under **Virtual media device**, click **Choose file**.
2. Select the file and click **Open**.
3. Click **Start** to start the session.

Server configuration
Learn about the tasks that are available from the **Server configuration** menu.

From this menu, you can choose from any of the following available tasks:

Network settings
Learn how to view and set common network, IPv4, and DNS settings.

To view network settings, select the **Network Interface** that you want to view. The **Hostname**, **MAC Address**, and **Default Gateway** are displayed under **Common settings**. The **DHCP setting**, **IPv4 IP addresses**, **Gateways**, and **Netmasks** are displayed under **IPv4 Settings**. Under **DNS settings**, all DNS servers are displayed.

To set network settings, complete the following steps:

1. Select the **Network Interface** that you want to set.
2. Edit the **Hostname**, **MAC Address**, or **Default Gateway** fields under **Common settings**.
3. Edit the **DHCP setting**, **IPv4 IP address**, **Gateway**, and **Netmask Prefix Length** under **IPv4 Settings**.
4. Click **Save Settings**.

Note: You can edit network settings only on firmware level OP920.01 or later.

To add an IPv4 address, complete the following steps:

1. Click **Add IPV4 address**.
2. Complete the **IPv4 IP address**, **Gateway**, and **Netmask Prefix Length** fields.
3. Click **Save Settings**.

To add a DNS address, complete the following steps:

1. Click **Add DNS Server**.
2. Enter the Internet Protocol (IP) address of the **DNS Server**.
3. Click **Save Settings**.

SNMP settings
Learn how to view and set the simple network management protocol (SNMP) with a hostname or Internet Protocol (IP) address and a port.

To set the SNMP, complete the following steps:

Note: Only SNMPv2 is supported.

1. Click **Add Manager**.
2. Enter the hostname or IP address and the port number.
3. Click **Save Settings**.

You can remove a manager by clicking the trash bin icon next to the manager that you want to remove.

Firmware
Learn how to manage the BMC and server firmware.

You can use the **BMC images** and **Server images** tables to manage firmware image files. The image file that is listed at the top, the image with the highest boot priority, is used the next time that the device is booted. You can change the boot order for the image file by clicking the arrow icons.

Learn about the different image states that are available:

- **Functional**: The running image on the device.
- **Active**: The image is available to boot from, but is not currently the running image. If the image is the top image in the relevant table, it becomes the functional image the next time the device is rebooted.
- **Activating**: The image is in the process of being activated and becomes either **Active** or **Failed**.
- **Failed**: The image failed to activate.
- **Ready**: The image is ready to be activated.
- **Invalid**: This image is an invalid image and cannot be activated.

The system firmware is a combination of the BMC firmware and the PNOR firmware. You must update both the BMC firmware and the PNOR firmware for the system to operate properly. If you update only one type of firmware, but do not update the other type of firmware, system errors might occur.

You can upload an image file from the workstation or from the Trivial File Transfer Protocol (TFTP) server. If you choose **Upload image file from workstation**, click **Choose a file** and specify the location of the image on the workstation storage device. Click **Upload** to upload the image file to the BMC server. If you choose **Download image file from TFTP server**, enter the TFTP server IP address in the **TFTP Server IP Address** field and the file name in the **File Name** field. Click **Download** to download the image file to the BMC server.

After you load the new image file to the BMC server, you can activate the image file to make it available for use. Locate the image in the correct image table, and then click **Activate** > **Continue**. For a BMC image, an option of **Activate Firmware File Without Rebooting BMC** or **Activate Firmware File and Automatically Reboot BMC** is available. If you select **Activate Firmware File Without Rebooting BMC**, the BMC must be rebooted by using the **Reboot BMC** option to make the image become the **Functional** image. If you select **Activate Firmware File and Automatically Reboot BMC**, the BMC automatically reboots after the image is activated and the new image becomes the **Functional** image.

For a server image, after the image is activated, the server must be rebooted (or powered on if the server is off) for the image to become active. The **Reboot** (or **Power on**) option can be accessed from the **Server power operations** menu.

Date and time settings
Learn how to set the date and time.

To automatically set the date and time, complete the following steps:

1. Select **Obtain automatically from a network time protocol (NTP) server**.
2. Click **Add new NTP server**.
3. Enter the NTP server address.
4. Click **Save setttings**.

To manually add the date and time, complete the following steps:

1. Select **Manually set date and time**.

2. Enter the date and time.
3. Change the **Time owner** to the following values:
   - **BMC**: the BMC owns the time and can set the time.
   - **Host**: the host owns the time and can set the time.
   - **Split**: the BMC and the host own separate time.

     Note: Split is the recommended value for **Time owner** for 7063-CR2.
   - **Both**: both the BMC and the host can set the time.

Access control
Learn about the tasks that are available from the **Access control** menu.

From this menu, you can choose from any of the following available tasks:

LDAP
Learn how to configure lightweight directory access protocol (LDAP) settings and manage role groups.

LDAP authentication

To enable LDAP authentication, complete the following steps:

1. Select the **Enable LDAP authentication** checkbox.

   Note: If you want to secure LDAP by using Secure Sockets Layer (SSL), select the **Secure LDAP using SSL** checkbox. You must have a certificate authority (CA) and LDAP certificate for this function.
2. Select the service type as **Open LDAP** or **Active directory**.
3. Complete the required fields.
4. Click **Save**.

Role groups

To add a new role group, complete the following steps:

1. Click **Add role group**.
2. Enter a name for the role group.
3. Set the privilege of the role group to **Administrator**, **Operator**, **User**, or **Callback**.
4. Click **Save**.

To remove a new role group, complete the following steps:

1. Select the checkbox next to the role group or groups that you want to remove from the table.
2. Click **Remove role groups**.
3. Click **Remove** in the pop-up window.

To modify the privilege of a role group, complete the following steps:

Note: LDAP authentication must be enabled to modify group roles.

1. Select the checkbox next to the role group or groups that you want to modify from the table.
2. Click the **Edit** icon.
3. Change the privilege of the role group to **Administrator**, **Operator**, **User**, or **Callback**.
4. Click **Save**.

Local users
Learn how to add or remove new users, modify user settings, manage user account policy settings, and view privilege role descriptions.

To add a new user, complete the following steps:

1. Click **Add user**.
2. Set the account status to either **Enabled** or **Disabled**.
3. Enter a new username.

   Note: The username cannot start with a number. No special characters are allowed except for an underscore.
4. Set the privilege of the user to **Administrator**, **Operator**, **ReadOnly**, or **NoAccess (Callback)**.
5. Enter the password of the user.

   Note: Initially, the password must be in the range of 8 - 20 characters in length. Passwords can be changed later to a length greater than 20 characters, but IPMI access is removed. The system might take up to 5 minutes for the new password to update on the BMC. If you have trouble accessing your account, wait for 5 minutes and try again.
6. Reenter the password for confirmation.
7. Click **Add user**.

To remove a user, complete the following steps:

1. Click the checkbox next to the user or users that you want to remove from the table.
2. Click **Remove**.
3. Click **Remove** again in the pop-up window.

You can modify user settings by selecting the user from the table and clicking the edit icon. From the **Modify user** window, you can update the following properties:

- Account status: set to **Enabled** or **Disabled**.
- Username: change the name of the user.
- Privilege: change the account privileges of the user to **Administrator**, **Operator**, **ReadOnly**, or **NoAccess (Callback)**.
- User password: update the password of the user.

You can modify user account policy settings by clicking **Account policy settings**. From the **Account policy settings** window, you can update the following properties:

- Maximum failed login attempts: change the number of allowed failed login attempts.
- User unlock method: set to **Automatic after timeout** or **manual**.
- Only non-root accounts can be locked out.

You can view privilege role descriptions by clicking **View privilege role descriptions**.

Table 31. Privilege role descriptions

Role	Privileges	Guidance
Administrator	Can configure the BMC and manage users and sessions. Operational control over BMC functions.	Use this role for the most trusted users.
Operator	Operational control over BMC functions.	Use this role for users who manage routine operations.
ReadOnly	Read-only access to BMC functions.	Use this role for users who need to monitor the BMC, but do not need to operate it.
NoAccess (Callback)	No access to BMC functions.	Use this role for users who do not need access to the web interface or REST APIs.

SSL certificates
Learn how to generate a certificate signing request (CSR), add new certificates, and replace existing certificates.

Generating a CSR

To generate a new CSR, complete the following steps:

1. Click **Generate CSR**.
2. Complete the required fields under the **General** section.
3. Under **Private key** > **Key Pair Algorithm**, select the algorithm as **EC** or **RSA**.
4. Click **Generate CSR**.

Certificates

To add a new certificate, complete the following steps:

1. Click **Add new certificate**.
2. Select the certificate type as **HTTPS Certificate**, **LDAP Certificate**, or **CA Certificate**.
3. Click **Choose file** to select the certificate.
4. Click **Open**.
5. Click **Save**.

To replace certificates, complete the following steps:

1. Select the certificate that you want to replace from the table.
2. Click the refresh icon.
3. Click **Choose file** to select the new certificate.
4. Click **Open**.
5. Click **Replace**.

Managing the system by using DMTF Redfish APIs
OpenBMC-based systems can be managed by using the DMTF Redfish APIs.

Overview

Redfish is a REST API used for platform management and is standardized by the Distributed Management Task Force, Inc. (http://www.dmtf.org/standards/redfish).

Redfish enables platform management tasks to be controlled by client scripts that are developed by using secure and modern programming paradigms.

The Redfish API enables provisioning of tunable parameters for better utilization of power.

IBM OpenBMC-based systems support DMTF Redfish API (DSP0266, version 1.7.0, published on 20 May 2019) for systems management.

A copy of the Redfish schema files that are in JSON format are published by DMTF (http://redfish.dmtf.org/schemas/v1/) and are packaged in the firmware image.

The schema files that are distributed in the chip enable proper functioning of the APIs in deployments that have no wide area network (WAN) connectivity.

Note: The Redfish API is enabled by default and the Redfish service cannot be enabled or disabled by the user.

Firmware levels

Redfish APIs are supported on OpenPOWER (OP) firmware level OP940, or later.

Communication prerequisites for Redfish on OpenBMC-based servers

Depending on the current firmware level and network deployment, complete the following prerequisite tasks:

- Upgrade the server firmware level to OP940, or later.
- Identify the IP address of the BMC.
- Install and run cURL (https://curl.haxx.se/) with the method, Uniform Resource Indicator (URI), and the request body as parameters to communicate with the Redfish service.
- Install Python on the client system (typically a Linux host).
- Optionally, install and run DMTF Redfishtool (https://github.com/DMTF/Redfishtool).

Interacting with the Redfish service

To interact with the Redfish service, complete the following steps.

1. Create an authenticated login session (POST method on the `/redfish/v1/SessionService/Sessions` resource).
2. Extract and save the following details:
   - Authentication token (found in the **X-Auth-Token** header of the response)
   - Session URI (found in the **Location** header of the response)
3. To read the properties of a resource, send a **GET** request with the **X-Auth-Token** header for the URI of the resource.
4. To set a property of a resource, send a **PATCH** request with the **X-Auth-Token** header for the URI of the resource, the property name, type, and value encoded as a JSON body.
5. Extract and parse the response from the Redfish service that contains the JSON body.

Redfish service home page URI

The Redfish service home page URI (also known as the service ROOT) can be accessed by retrieving the URI: `https://<`*`ip:port`*`>/redfish/v1`. The response to this URI is a high-level site map that enables a traversal of the Redfish service by using a hypermedia API paradigm.

Interpreting the data returned by the Redfish service

The format and structure of the data is defined in the schema files. Schema files are JSON files that describe the data that is sent by the Redfish service. You can use the schema files to understand the data that is sent by the Redfish service and to validate the response that is sent by the Redfish service.

Location of the schema files

DMTF publishes the schema files for the standard data that is used in Redfish.

The Redfish schema files in JSON format are hosted in the DMTF schema repository at http://redfish.dmtf.org/schemas/v1/

Supported schema files

The following schema files are supported for OpenBMC-based systems:

- Account
- AccountCollection
- AccountService
- Certificate
- CertificateCollection

- CertificateLocations
- CertificateService
- Chassis
- ChassisCollection
- ComputerSystem
- ComputerSystemCollection
- EthernetInterface
- EthernetInterfaceCollection
- LogEntry
- LogService
- LogServiceCollection
- Manager
- ManagerCollection
- ManagerNetworkProtocol
- Memory
- MemoryCollection
- Processor
- ProcessorCollection
- Role
- RoleCollection
- ServiceRoot
- Session
- SessionCollection
- SessionService
- SoftwareInventory
- SoftwareInventoryCollection
- ThermalPower
- UpdateService

Accessing the common system management functions on the Redfish service by using cURL command

The following examples show the client URL (cURL) commands that can be used to access the common functions that are supported by the OpenBMC Redfish APIs:

Note: In all cURL commands, *${BMC}* is the IP address of the BMC.

- To view major collections, run the following commands:
  - Chassis collection:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Chassis
    ```

  - Manager collection:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Managers
    ```

  - System collection:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Systems
    ```

- To view the chassis, manager, and system resources, run the following commands:
  - Chassis resource:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Chassis/chassis
    ```

  - Manager resource:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Managers/bmc
    ```

  - System resource:

    ```
    curl -u root:0penBmc -k -s  https://${BMC}/redfish/v1/Systems/system
    ```

- To perform host power control operations, run the following commands:
  - Host power on:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "On"}'
    ```

  - Host soft power off:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "GracefulShutdown"}'
    ```

  - Host hard power off:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "ForceOff"}'
    ```

  - Restart host:

    ```
    -X POST https://${BMC}/redfish/v1/Systems/system/Actions/ComputerSystem.Reset -d
    '{"ResetType": "GracefulRestart"}'
    ```

- To view the host power control resource, run the following command:

  ```
  curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/Actions/
  ```

- To view the log resource, run the following command:

  ```
  curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/LogServices/EventLog/
  Entries
  ```

- To view sensor resources, run the following commands:
  - Power resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/chassis/Power
    ```

  - Thermal resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/chassis/Thermal
    ```

  - Sensor resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/chassis/Sensors
    ```

- To view inventory resources, run the following commands:
  - Memory resource:

    ```
    curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/Memory
    ```

  - Processor resource:

```
curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Systems/system/Processors
```

  - Power supply 0 resource:

```
curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/powersupply0
```

  - Power supply 1 resource:

```
curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/powersupply1
```

  - Motherboard resource:

```
curl -u root:0penBmc -k -s https://${BMC}/redfish/v1/Chassis/motherboard
```

- To update the firmware, run the following commands:
  - By using an image file from your system:

```
curl -u root:0penBmc -curl k -s  -H "Content-Type: application/octet-stream" -X POST -T
<image file path> https://${BMC}/redfish/v1/UpdateService
```

  - By using a Trivial File Transfer Protocol (TFTP) server:

```
curl -u root:0penBmc -k -s -d '{"ImageURI":"<TFTP IP Address>/<File name on
TFTP server>","TransferProtocol":"TFTP"}' -X POST https://${BMC}/redfish/v1/UpdateService/
Actions/UpdateService.SimpleUpdate
```

- To create a new local account, run the following command:
  -

```
curl -X POST https://${BMC}/redfish/v1/AccountService/Accounts/ -d '{"UserName": "admin",
"Password": "NEWPASSWORD", "RoleId": "Administrator"}'
```

Where `admin` is the name of the user that you want to create, `NEWPASSWORD` is the new password, and `RoleId` maps to the privilege role.

- To change the account password, run the following command:
  -

```
curl -X POST https://${BMC}/redfish/v1/AccountService/Accounts/root -d '{"Password":
"NEWPASSWORD"}'
```

Where `root` is the account name or user ID and `NEWPASSWORD` is the new password.

For more information about selecting a username, password, or role, see "Local users" on page 57.

Configure BMC connectivity (7063-CR2)
You can configure or view the network settings on the BMC for the management console.

Notes:

- This task applies only to the 7063-CR2. This connection is required to access the baseboard management controller (BMC) on the HMC.
- The settings in this task are applicable only to the dedicated IPMI or BMC network port.

To configure the BMC connection, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change BMC/IPMI network settings**.
3. Select the connection mode (**DHCP** or **Static**).

   If you select **Static** mode, complete the following addresses:

   - **IP address**
   - **Subnet mask**
   - **Gateway**
4. Click **OK**.

You can also configure the BMC network connection by using the Petitboot bootloader interface. For more information, see Configuring the firmware IP address.

Setting the IPv4 address
Learn how to set your IPv4 address on the HMC.

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **Basic Settings** tab.
6. Select an IPv4 address.
7. If you selected to specify an IP address, enter the TCP/IP interface address and the TCP/IP interface network mask.
8. Click **OK**.

Setting the IPv6 address
Learn how to set your IPv6 address on the HMC.

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **IPv6 Settings** tab.
6. Select an **Autoconfig** option or add a static IP address.
7. If you added an IP address, enter the IPv6 address and the prefix length and click **OK**.
8. Click **OK**.

Using only IPv6 addresses
Learn how to configure the HMC so that it uses only IPv6 addresses.

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Select **No IPv4 address**.
6. Click the **IPv6 Settings** tab.
7. Select **Use DHCPv6 to configure IP settings** or add static IP addresses, then click **OK**.

What to do next

After you click **OK**, you must restart your HMC for these changes to take effect.

Changing HMC firewall settings

In an open network, a firewall is used to control outside access to your company network. The HMC also has a firewall on each of its Ethernet adapters. To control the HMC remotely or give remote access to others, modify the firewall settings of the Ethernet adapter on the HMC that is connected to your open network.

About this task

To configure a firewall, use the following steps:

Procedure

1. In the navigation area, click the **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**.
3. Click the **LAN Adapters** tab.
4. Select the LAN adapter that you want to work with and click **Details**.
5. Click the **Firewall** tab.
6. Using one of the following methods, you can allow any IP address by using a particular application through the firewall, or you can specify one or more IP addresses:
   - Allow any IP address by using a particular application through the firewall:
     a. From the top box, highlight the application.
     b. Click **Allow Incoming**. The application displays in the bottom box to signify that it is selected.
   - Specify which IP addresses to allow through the firewall:
     a. From the top box, highlight an application.
     b. Click **Allow Incoming by IP Address**.
     c. On the Hosts Allowed window, enter the IP address and the network mask.
     d. Click **Add** and click **OK**.
7. Click **OK**.

   Notes:

   - For more information about enabling remote restricted shell access, see "Enabling remote restricted shell access" on page 65.
   - For more information about enabling remote web access, see "Enabling remote web access" on page 66.

Enabling remote restricted shell access
You can enable remote restricted shell access when you configure a firewall.

About this task

To enable remote restricted shell access, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security** > **Systems and Console Security**.
2. Click **Enable Remote Command Execution**.
3. Select **Enable remote command execution using the ssh facility** and then click **OK**.

What to do next

Now remote restricted shell access is enabled.

Enabling remote web access

You can enable remote web access to your Hardware Management Console (HMC).

About this task

To enable remote web access, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security** > **Systems and Console Security**.
2. Click **Enable Remote Command Execution**.
3. Select **Enable** and then click **OK**.

What to do next

Now remote web access is enabled.

Configuring a routing entry as the default gateway

Learn how to configure a routing entry as the default gateway. This task is available when you are using an open network.

Before you begin

To configure a routing entry as the default gateway, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Customize Network Settings window opens.
3. Click the **Routing** tab.
4. In the Default gateway information section, enter the gateway address and gateway device of the routing entry you want to set as the default gateway.
5. Click **OK**.

Configuring domain name services

If you plan to set up an open network, configure domain name services.

About this task

If you plan to set up an open network, configure domain name services. Domain Name System (DNS) is a distributed database system for managing host names and their associated Internet Protocol (IP) addresses. Configuring domain name services includes enabling DNS and specifying the domain suffix search order.

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Change Network Settings window opens.
3. Click the **Name Services** tab.
4. Select **DNS enabled** to enable DNS.
5. Specify the DNS server and domain suffix search order and click **Add**.
6. Click **OK**.

Configuring domain suffixes

The list of domain suffixes is used to resolve an IP address that starts with the first entry in the list.

About this task

The domain suffix is a string that is appended to a host name that is used to help resolve its IP address. For example, a host name of `myname` might not be resolved. However, if the string `myloc.mycompany.com` is an element in the domain suffix table, then an attempt is made to resolve `myname.mloc.mycompany.com`.

To configure a domain suffix entry, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change network settings**. The Customize Network Settings window opens.
3. Click the **Name Services** tab.
4. Enter a string to be used as a domain suffix entry.
5. Click **Add** to add it to the list.

Configuring the HMC so that it uses LDAP remote authentication

You can configure your Hardware Management Console (HMC) so that it uses LDAP (Lightweight Directory Access Protocol) remote authentication.

Before you begin

When a user logs in to the HMC, authentication is first performed against a local password file. If a local password file is not found, the HMC can contact a remote LDAP server for authentication. You must configure your HMC so that it uses LDAP remote authentication.

Note: Before you configure the HMC so that it uses LDAP authentication, you must ensure that a working network connection exists between the HMC and the LDAP servers.

About this task

To configure your HMC so that it uses LDAP authentication, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security**, and then select **Systems and Console Security**.
2. In the content pane, select **Manage LDAP**. The LDAP Server Definition window opens.
3. Select **Enable LDAP**.
4. Define an LDAP server to use for authentication.
5. Define the LDAP attribute that is used to identify the user that is being authenticated. The default is **uid**, but you can use your own attributes.
6. Define the distinguished name tree, also known as the search base, for the LDAP server.
7. Click **OK**.
8. If a user wants to use LDAP authentication, the user must configure their profile so that it uses LDAP remote authentication instead of local authentication.

Configuring the HMC so that it uses Key Distribution Center servers for Kerberos remote authentication

You can configure the HMC so that it uses Key Distribution Center (KDC) servers for Kerberos remote authentication.

Before you begin

When a user logs in to the HMC, authentication is first verifies against a local password file. If a local password file is not found, the HMC can contact a remote Kerberos server for authentication. You must configure your HMC so that it uses Kerberos remote authentication.

Note: Before you configure the HMC so that it uses KDC servers for Kerberos remote authentication, you must ensure that a working network connection exists between the HMC and the KDC servers.

About this task

To configure the HMC so that it uses KDC servers for Kerberos remote authentication, complete the following steps:

Procedure

1. Enable the Network Time Protocol (NTP) service on the HMC and set the HMC and the KDC servers to synchronize time with the same NTP server. To enable the NTP service on the HMC, complete the following steps:
   a) In the navigation area, click **Console Management**, and then select **Console Settings**.
   b) In the content pane, select **Change Date and Time**.
   c) Select the **NTP Configuration** tab.
   d) Select **Enable NTP service on this HMC**.
   e) Click **OK**.
2. Configure each remote HMC user's profile so that it uses Kerberos remote authentication instead of local authentication.
3. Optionally, you can import a service-key file into this HMC. The service-key file contains the host principal that identifies the HMC to the KDC server. Service-key files are also known as *keytabs*. To import a service-key file into this HMC, complete the following steps:
   a) In the navigation area, click **Users and Security**, and then select **Systems and Console Security**.
   b) In the content pane, select **Manage KDC**.
   c) Select **Actions > Import Service Key**. The Import Service Key window opens.
   d) Type the location of the service key file.
   e) Click **OK**.
4. Add a new KDC server to this HMC. To add a new KDC server to this HMC, complete the following steps:
   a) In the navigation area, click **Users and Security**, and then select **Systems and Console Security**.
   b) In the content pane, select **Manage KDC**.
   c) Select **Actions > Add KDC Server**. The Import Service Key window opens.
   d) Type the realm and the host name or IP address of the KDC server.
   e) Click **OK**.

Configuring the local console to report errors to service and support

Configure this HMC so that it can call-home errors by using LAN connectivity.

Configuring the HMC so that it can connect to service and support by using the call-home setup wizard
Configure the HMC so that it is a call-home server by using the call-home wizard.

Before you begin

This procedure describes how to configure the HMC as a call-home server by using direct (LAN-based) and indirect (SSL) connections to the internet.

Before you begin this task, ensure that:

- The network administrator verifies that connectivity is allowed.
- If you are configuring internet support through a proxy server, you must also have the following information:
  - The IP address and port of the proxy server
  - The proxy authentication information
- The adapter that is designated as **eth1** (the one that is designated as an open network) is used. For more information, see "Choosing network settings on the HMC" on page 15.
- An Ethernet cable physically connects the HMC to the LAN.

To configure the HMC so that it is a call-home server by using the call-home wizard, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Call-Home Setup Wizard**. The Connectivity and Call-Home Servers wizard opens. Follow the instructions in the wizard to configure call-home.

Configuring the local console to report errors to service and support
Configure this HMC so that it can call-home errors by using LAN connectivity.

Configuring an HMC to contact service and support by using LAN-based internet and SSL
Describes how to configure the HMC as a call-home server by using direct (LAN-based) and indirect (SSL) connections to the internet.

Before you begin

Before you begin this task, ensure that:

- The network administrator verifies that connectivity is allowed..
- Customer contact information is configured. Verify the contact information by going to the HMC interface and clicking **Serviceability>Service Management > Manage Customer Information**.
- If you are configuring internet support through a proxy server, you must also have the following information:
  - The IP address and port of the proxy server
  - The proxy authentication information
- You need at least one open network interface configured. For more information, see "Private and open networks in the HMC environment" on page 17.
- An Ethernet cable physically connects the HMC to the LAN.

About this task

To configure the HMC as a Call Home server by using LAN-based internet and SSL, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the Connectivity section, click **Manage Outbound Connectivity**. The Call-home Server Consoles window opens.
3. Click **Configure.**
4. In the Outbound Connectivity Settings window, check **Enable local system as call-home server**.
5. Accept the agreement.
6. In the Outbound Connectivity Settings window, select the **Internet** page.
7. Check the **Allow an existing internet connections for service** box.
8. If you are using an SSL proxy, check the **Use SSL proxy** box.
9. If you are using an SSL proxy, complete the proxy's address and port. Obtain this information from the network administrator.
10. If you checked **Use SSL proxy** and the proxy requires user ID and password authentication, check the **Authenticate with the SSL proxy** box. Type the user ID and password. Obtain the user ID and password from the network administrator.
11. Select the **Protocol to Internet** you want to use.
12. On the **Internet** page, click **Test**.
13. In the Test internet window, click **Start**.
14. Verify that the test completes successfully.
15. In the Test internet window, click **Cancel**.
16. In the Outbound Connectivity Settings window, click **OK**.

Choosing existing call-home servers to connect to service and support for this HMC
Choose existing Hardware Management Console (HMC) call-home servers that are recognized or discovered by the HMC to report errors.

Before you begin

Discovered HMCs are HMCs that are enabled as call-home servers and are either on the same subnet or manage the same managed system as this HMC.

To choose a discovered HMC to call home when the HMC reports errors, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Outbound Connectivity**. The Call-Home Server Consoles window opens.
3. Click **Use discovered call-home server consoles**. The HMC displays the IP address or host name of the HMCs configured for call-home.
4. Click **OK**.

Results

You can also manually add existing HMC call-home servers that are on a different subnet. Select the IP address or host name of the HMC that is configured for call home and click **Add** and then click **OK**.

Verifying that your connection to service and support is working
Test problem reporting to ensure that connection to service and support is working.

About this task

To verify that your call-home configuration is working, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Create Event**.
3. Select **Test Automatic problem Reporting** and type a comment.
4. Click **Request Service**. Wait a few minutes for the request to be sent.
5. In the Service Management window, select **Manage Events**.
6. Select **All open problems**.
7. Verify that a PMH event and number is assigned to the problem number you opened.
8. Select that event and click **Close**.
9. On the **Close** window, type your name and a brief comment.

Authorizing users to view collected system data
You must authorize users to view data about your systems.

Before you begin

Before you authorize users to view collected system data, you must obtain an IBM ID. For more information about obtaining an IBM ID, see “Preinstallation configuration worksheet for the HMC” on page 24.

About this task

To authorize users to view collected system data, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, select **Authorize User**.
3. Enter your IBM ID.
4. Click **OK**.

Transmitting service information
You can transmit information to your service provider immediately, or you can schedule the information to be sent regularly.

Before you begin

IBM provides personalized web functions that use information that is collected by IBM Electronic Service Agent. To use these functions, you must first register on the IBM Registration website at http://www.ibm.com/account/profile. To authorize users to use the Electronic Service Agent information to personalize the web functions, see “Authorizing users to view collected system data” on page 71. For more information about the benefits of registering an IBM ID with your systems, see http://www.ibm.com/support/electronic.

Note: You must transmit service provider information as soon as the HMC is installed and configured for use.

About this task

To transmit service information, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Transmit Service Information.**
3. Complete the tasks in the **Transmit Service Information** window, and click **OK**.

Configuring the Events Manager for Call Home

Learn how to configure the Events Manager for Call Home task. You can monitor and approve any data that is being transmitted from an HMC to IBM through this task.

The Events Manager for Call Home mode (enabled or disabled) is set by using the HMC command line interface. Enabling the Events Manager for Call Home task blocks the HMC from automatically calling home events as they occur. To prevent events that are called home without approval, all HMCs running in this environment must have the Events Manager for Call Home enabled.

To enable or disable the Events Manager for Call Home task, run the following command:

```
chhmc -c emch

-s {enable | disable}

[--callhome {enable | disable}]

[--help]
```

Note: Enabling the Events Manager for Call Home task holds call home events until they are approved for the call home task. If you disable the Events Manager for Call Home task, it does not automatically enable the call home feature. This setup prevents any unintended call home of data back to IBM. Choose from the following command options to set up the required configuration:

- To enable the Events Manager for Call Home task: **`chhmc -c emch -s enable`**
- To disable the Events Manager for Call Home task and to re-enable automatic call home: **`chhmc -c emch -s disable --callhome enable`**
- To disable the Events Manager for Call Home task and not re-enable automatic call home: **`chhmc -c emch -s disable --callhome disable`**

Ensure that the HMC can communicate with other HMCs deployed in this environment. The Events Manager for Call Home has a test connection function when an HMC is registered.

You can register the HMC with the Events Manager for Call Home. After you register the HMC, the events manager queries the registered HMC for any events that are waiting to be called home to IBM. The Events Manager shows what data is being sent back to IBM and approves these events. After approval, the Event Manager notifies the registered HMC that it can proceed with the call home operation.

The Events Manager for Call Home task can be run from any HMC or from multiple HMCs. To register a management console with the Events Manager for Call Home task, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Events Manager for Call Home**.
2. From the **Events Manager for Call Home** pane, click **Manage Consoles**.
3. From the **Manage Registered Consoles** window, click **Add Console** to enter information to register a management console with the Events Manager for Call Home task.
4. Click **OK** to commit the changes to the list of registered management console.

Note: The Events Manager for Call Home can be used with the event manager mode disabled. You can still register the HMC and view events in the events manager, but Events Manager does not control when the events are called home.

Setting passwords for the managed system

You must set passwords for both your server and Advanced System Management (ASM). Read more about how to use the HMC interface to set these passwords.

Before you begin

If you received the message `Authentication Pending`, the HMC prompts you to set the passwords for the managed system.

About this task

If you did not receive the message Authentication Pending, complete the following steps to set the passwords for the managed system.

Updating your server password

Before you begin

To update your server password, complete the following steps:

Procedure

1. In the navigation area, select the managed system and click **Users and Security**, and then select **Users and Roles**.
2. Click **Change Password**. The Update Password window opens.
3. Type the required information and click **OK**.

Updating your Advanced System Management (ASM) general password

Before you begin

Note: The default password for the general user ID is `general`, and the default password for the administrator ID is `admin`.

To update your ASM general password, complete the following steps:

Procedure

1. In the navigation area of the HMC, select the managed system.
2. In the Tasks area, click **Operations**.
3. Click **Advanced System Management (ASM)**. The Launch ASM Interface window opens.
4. Select a Service Processor IP Address and click **OK**. The ASM interface opens.
5. On the ASMI Welcome pane, specify your user ID and password, and click **Log In**.
6. In the navigation area, expand **Login Profile**.
7. Select **Change Password**.
8. Specify the required information, and click **Continue**.

Resetting the Advanced System Management (ASM) administrator password

Before you begin

To reset the administrator password, contact an authorized service provider.

Testing the connection between the HMC and the managed system

Learn how to verify that you are properly connected to the network.

About this task

To test the network connectivity, you must be a member of one of the following roles:

- Super administrator
- Service representative

To test the connection between the HMC and the managed system, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Settings**.

2. In the content pane, click **Test Network Connectivity**.
3. In the Ping tab, type the host name or IP address of any system to which you want to connect. To test an open network, type the gateway. Click **Ping**.

Results

If you have not created any logical partitions, you cannot ping the addresses. You can use the HMC to create logical partitions on your server. For more information, see Logical partitioning.

To understand how the HMC can be used in a network, see "HMC network connections" on page 15.

For more information about configuring the HMC to connect to a network, see "Configuring the HMC by using the menus " on page 31.

Postconfiguration steps

After you install and configure the HMC, back up HMC data as necessary.

Backing up management console data

This task backs up (or archives) the data that is stored on your HMC hard disk that is critical to support HMC operations.

Before you begin

Your remote system must have Network File System (NFS) or Secure Shell (ssh) configured, and this network must be accessible from the HMC. To complete this task, you must shut down and reboot the HMC. Use only the HMC to perform these tasks.

About this task

To back up the HMC hard disk drive to a remote system, you must be a member of one of the following roles:

- Super administrator
- Operator
- Service representative

Back up the HMC data after changes have been made to the HMC or information associated with logical partitions.

The HMC data stored on the HMC hard drive can be saved to a DVD-RAM on a local system, a remote system mounted to the HMC file system (such as NFS), or sent to a remote site using File Transfer Protocol (FTP).

Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

Using the HMC, you can back up all important data, such as the following:

- User-preference files
- User information
- HMC platform-configuration files
- HMC log files
- HMC updates through Install Corrective Service.

To back up the HMC hard drive to a remote system, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Backup Management Console Data**.

3. From the **Backup Management Console Data** window, select the archive option you want to perform.
4. Click **Next**, then follow the appropriate instructions depending on the option you chose.
5. Click **OK** to continue with the backup process.

Updating, upgrading, and migrating your HMC machine code

Updates and upgrades are periodically released for the HMC to add new functionality and to improve existing features. Learn more about the differences between updating, upgrading, and migrating your HMC machine code. Also learn how to perform an HMC machine code update, upgrade, or migration.

When you are finished with each of these tasks, the HMC reboots but the partitions do not.

Updating HMC code

Applies maintenance to an existing HMC level

Does not require that you perform the **Save upgrade data** task

Upgrading HMC code

Replaces HMC software with a new release or fix level of the same program

Requires that you boot from recovery media

Migrating HMC code

Moves HMC data from one HMC version to another

A migration is a type of upgrade.

Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

Determining your HMC machine code version and release

Find out how to view the HMC machine code version and release.

About this task

The level of machine code on the HMC determines the available features, including concurrent server firmware maintenance and enhancements to upgrading to a new release.

To view the HMC machine code version and release, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the **Current HMC Driver Information** heading, including: the HMC version, release, maintenance level, build level, and base versions.

Obtaining and applying machine code updates for the HMC with an Internet connection

Learn how to obtain machine code updates for the HMC when the HMC has an Internet connection.

About this task

To obtain machine code updates for the HMC, complete all steps.

Step 1. Ensure that you have an Internet connection

About this task

To download updates from the service and support system or website to your HMC or server, you must have one of the following connections:

- SSL connectivity with or without a SSL proxy
- Internet VPN

To ensure that you have an Internet connection, do the following:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, **Manage Outbound Connectivity**.
3. Select the tab for the type of outbound connectivity that you chose for your HMC (Internet VPN or SSL connectivity).

   Note: If a connection to service and support does not exist, set up the service connection before proceeding with this procedure. For instructions on how to set up a connection to service and support, see Setting up your server to connect to IBM service and support.
4. Click **Test**.
5. Verify that the test completes successfully.

   If the test is not successful, troubleshoot your connectivity and correct the problem before proceeding with this procedure. Alternatively, you can obtain the update on DVD.

   Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.
6. Continue with "Step 2. View the existing HMC machine code level" on page 76.

Step 2. View the existing HMC machine code level

About this task

To view the existing HMC machine code level, complete the following steps:

Procedure

1. In the navigation area, click the **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Continue with "Step 3. View the available HMC machine code levels" on page 76.

Step 3. View the available HMC machine code levels

About this task

To view the available HMC machine code levels, complete the following steps:

To find out if there are new HMC machine code updates available, contact service and support.

Procedure

1. From a computer or server with an Internet connection, go to http://www.ibm.com/eserver/support/fixes.
2. Select the appropriate family in the Product family list.
3. Select **Hardware Management Console** in the Product or fix type list.
4. Click **Continue**.

   The Hardware Management Console site is displayed.
5. Scroll down to your HMC Version level to view available HMC levels.

Note: If you prefer, you can contact service and support.

6. Continue with "Step 4. Apply the HMC machine code update" on page 77.

Step 4. Apply the HMC machine code update

About this task

To apply the HMC machine code update, complete the following steps:

Procedure

1. Before you install updates to the HMC machine code, back up critical console information on your HMC.

   For instructions, see "Backing up management console data" on page 74. Then continue with the next step.
2. In the navigation area, click **Console Management**, and then select **Console Management**.
3. In the content pane, click **Update the Hardware Management Console**. The Install Corrective Service Wizard opens.
4. Follow the instructions in the Wizard to install the update.
5. Shut down and then restart the HMC for the update to take effect.
6. Click **Log on and launch the Hardware Management Console web application**.
7. Log in to the HMC interface.

Step 5. Verify that the HMC machine code update installed successfully

About this task

To verify that the HMC machine code update installed correctly, complete the following steps:

Procedure

1. In the navigation area, click the **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Verify that the version and release match the update that you installed.
5. If the level of code displayed is not the level that you installed, perform the following steps:

   a. Select the network connection on the HMC.

   b. Retry the firmware update using a different repository.

   c. If the problem persists, contact your next level of support.

Obtaining and applying machine code updates for the HMC using DVD or an FTP server

Learn how to obtain machine code updates for the Hardware Management Console (HMC) by using DVD or an FTP server.

About this task

To obtain HMC machine code updates, complete all steps.

Note: For HMC model 7063-CR1, you can connect an external USB DVD drive.

Step 1. View the existing HMC machine code level

Before you begin

To view the existing HMC machine code level, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Continue with "Step 2. View the available HMC machine code levels" on page 78.

Step 2. View the available HMC machine code levels

Before you begin

To view the available HMC machine code levels, complete the following steps:

About this task

To find out if there are new HMC machine code updates available, contact IBM service and support.

Procedure

1. From a computer or server with an internet connection, go to the Fix Central website.
2. Scroll down to your HMC Version level to view available HMC levels.

   Note: If you prefer, you can contact IBM service and support.
3. Continue with "Step 3. Obtain the HMC machine code update" on page 78.

Step 3. Obtain the HMC machine code update

Before you begin

To obtain the HMC machine code update, complete the following steps:

About this task

To order the HMC machine code update on removable media, contact service and support.

You can order the HMC machine code update through the Fix Central website, contact service and support, or download it to an FTP server.

Ordering the HMC machine code update through the Fix Central website

1. From a computer or server with an Internet connection, go to the Fix Central website.
2. Under Supported HMC products, select the latest HMC level.
3. Scroll down to the File names / Package area and locate the update that you want to order.
4. In the Order column, select **Go**.
5. Click **Continue** to sign in with your IBM ID.
6. Follow the on-screen prompts to submit your order.

Downloading the HMC machine code update to removable media

1. From a computer or server with an Internet connection, go to Fix Central website.

2. Under Supported HMC products, select the latest HMC level.
3. Scroll down to the File names / Package area and locate the update that you want to download.
4. Click the update that you want to download.
5. Accept the license agreement, and save the update to your removable media.

What to do next

When you are finished, continue with "Step 4. Apply the HMC machine code update" on page 79.

Step 4. Apply the HMC machine code update

Before you begin

To apply the HMC machine code update, complete the following steps:

Procedure

1. Before you install updates to the HMC machine code, back up HMC data. For more information, see "Backing up management console data" on page 74.
2. If you obtained or created the update on DVD-RAM, insert it into the DVD drive on the HMC. If you obtained or created the update on a USB memory device, insert the memory device.
3. Before you install updates to the HMC machine code, back up critical console information on your HMC.

   For instructions, see "Backing up management console data" on page 74. Then continue with the next step.
4. In the navigation area, click **Console Management**, and then select **Console Management**.
5. In the content pane, click **Update the Hardware Management Console**. The Install Corrective Service Wizard opens.
6. Follow the instructions in the Wizard to install the update.
7. Shut down, restart, and log back in to the HMC for the update to take effect.
8. Continue with "Step 5. Verify that the HMC machine code update installed successfully" on page 79.

Step 5. Verify that the HMC machine code update installed successfully

Before you begin

To verify that the HMC machine code update installed successfully, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver Information heading, including: the HMC version, release, maintenance level, build level, and base versions.
4. Verify that the version and release match the update that you installed.
5. If the level of code that is displayed is not the level that you installed, perform the following steps:
   a. Retry the machine code update. If you created a DVD for this procedure, use a new media.
   b. If the problem persists, contact your next level of support.

Upgrading your HMC software

Learn how to upgrade the software on an HMC from one release to the next while you maintain your HMC configuration data.

About this task

To upgrade the machine code on an HMC, complete all steps.

Note: For HMC models 7063-CR1 and 7063-CR2, you can connect an external USB DVD drive.

Step 1. Obtain the upgrade

About this task

To order the HMC machine code upgrade on DVD-RAM, contact service and support.

You can order the HMC machine code upgrade through the Fix Central website.

To obtain the upgrade through the Fix Central website, complete the following steps:

Procedure

1. From a computer or server with an internet connection, go to the Hardware Management Console website at http://www-933.ibm.com/support/fixcentral/.
2. Click **Continue**.

   The Hardware Management Console site is displayed.
3. Navigate to the HMC version you want to upgrade to.
4. Locate the download and ordering section.

   Note: If you do not have access to the internet, contact IBM service and support to order the upgrade on DVD.
5. Follow the on-screen prompts to submit your order.
6. After you have the upgrade, continue with "Step 2. View the existing HMC machine code level" on page 80.

Step 2. View the existing HMC machine code level

About this task

To determine the existing level of machine code on an HMC, follow these steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Management**. In the navigation area, click **Updates**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver information heading, including the HMC version, release, maintenance level, build level, and base versions.
4. Continue with "Step 3. Back up the managed system's profile data" on page 80.

Step 3. Back up the managed system's profile data

About this task

To back up the managed system's profile data, complete the following steps:

Procedure

1. Select the system that you want to save the profile data.
2. Click **Actions** > **View All Actions** > **Legacy** > **Manage Partition Data** > **Backup**.
3. Type a backup file name and record this information.
4. Click **OK**.
5. Repeat these steps for each system.
6. Continue with "Step 4. Back up HMC data" on page 81.

Step 4. Back up HMC data

About this task

Back up HMC data before you install a new version of HMC software so that previous levels can be restored in the event of a problem while you upgrade the software. Do not use this critical console data after a successful upgrade to a new version of the HMC software.

Note: To back up to removable media, you need to have that media available.

To back up HMC data, complete the following steps:

Procedure

1. If you plan to back up to media, perform the following steps to format the media:
   a. Insert the media into the drive.
   b. In the navigation area, click **Serviceability**, and then select **Service Management**.
   c. In the content pane, click **Format Media**.
   d. Select the media type.
   e. Select the format type.
   f. Click **OK**.
2. In the navigation area, click **Console Management**, and then select **Console Management**.
3. In the content pane, click **Backup Management Console Data**.

   The **Backup Management Console Data** window opens.
4. Select an archive option.

   You can back up to media on a local system, a remote system that is mounted to the HMC file system (for example, NFS), or send the backup to a remote site by using File Transfer Protocol (FTP).

   - To back up to a local system, choose **Back up to media on local system** and follow the instructions.
   - To back up to a mounted remote system, choose **Back up to mounted remote system** and follow the instructions.
   - To back up to a remote FTP site, choose **Send back up critical data to remote site** and follow the instructions.
5. Continue with "Step 5. Record the current HMC configuration information" on page 81.

Step 5. Record the current HMC configuration information

About this task

Before you upgrade to a new version of HMC software, as a precautionary measure, record HMC configuration information.

To record the current HMC configuration, complete the following steps:

Procedure

1. Select a managed system or any partitions that you want to record HMC configuration information.
2. From the menu pod, select **Actions** > **Schedule Operations**.

   All scheduled operations for the target that you selected are displayed.
3. Select **Sort** > **By Object**.
4. Select each object and record the following details:
   - Object Name
   - Schedule date
   - Operation Time (displayed in 24-hour format)
   - Repetitive (if Yes, complete the following steps):
     a. Select **View** > **Schedule Details**.
     b. Record the interval information.
     c. Close the scheduled operations window.
     d. Repeat for each scheduled operation.
5. Close the **Customize Scheduled Operations** window.
6. Continue with "Step 6. Record remote command status" on page 82.

Step 6. Record remote command status

About this task

To record remote command status, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security**, and then select **Systems and Console Security**.
2. In the content pane, click **Enable Remote Command Execution**.
3. Record whether the **Enable remote command execution using the ssh facility** check box is selected.
4. Click **Cancel**.
5. Continue with "Step 7. Save upgrade data" on page 82.

Step 7. Save upgrade data

About this task

You can save the current HMC configuration in a designated disk partition on the HMC or to local media. Save upgrade data only immediately before you upgrade your HMC software to a new release. You can restore the HMC configuration settings after you upgrade.

Note: Only one level of backup data is allowed. Each time that you save upgrade data, the previous level is overwritten.

To save upgrade data, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Save Upgrade Data**. The **Save Upgrade Data** wizard opens.
3. Select the media on which you want to save the upgrade data. If you choose to save to removable media, insert the media now. Click **Next**.
4. Click **Finish**.
5. Wait for the task to complete.

If the Save Upgrade Data task fails, contact your next level of support before proceeding.

Note: If the save upgrade data task fails, do not continue the upgrade process.

6. Click **OK**.
7. Continue with "Step 8. Upgrade the HMC software" on page 83.

Step 8. Upgrade the HMC software

About this task

To upgrade the HMC software, restart the system with the removable media in the DVD drive.

Procedure

1. Insert the HMC Product Installation media into the DVD drive.
2. In the navigation area, click **Console Management**, and then select **Console Management**.
3. In the content pane, select **Shutdown or Restart the Managment Console**.
4. Ensure **Restart the HMC** is selected.
5. Click **OK**.

   The HMC restarts and system information scrolls on the window.
6. Select **Upgrade** and click **Next**.
7. Choose from the following options:
   - If you saved the upgrade data during the previous task, continue with the next step.
   - If you did not save the upgrade data previously in this procedure, you must save the upgrade data now before you continue.
8. Select **Upgrade from media** and click **Next**.
9. Confirm the settings and click **Finish**.
10. Follow the prompts.

    Note:

    - If the screen goes blank, press the space bar to view the information.
    - The first DVD can take approximately 20 minutes to install.
11. At the login prompt, log in using your user ID and password.

    The HMC code installation is complete.
12. Continue with "Step 9. Verify that the HMC machine code upgrade installed successfully" on page 83.

Step 9. Verify that the HMC machine code upgrade installed successfully

About this task

To verify that the HMC upgrade is installed successfully, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**.
3. In the new window, view and record the information that appears under the Current HMC Driver information heading, including the HMC version, release, maintenance level, build level, and base versions.
4. Verify that the version and release match the update that you installed.
5. If the level of code that is displayed is not the level that you installed, retry the upgrade task by using a new DVD. If the problem persists, contact your next level of support.

Upgrading HMC from remote location by using network upgrade images

Learn how to upgrade the software on an HMC from a remote location by using network upgrade images.

About this task

Learn how to upgrade the software on an HMC from a remote location by using network upgrade images.

Procedure

1. From a computer or server with an internet connection, go to the Hardware Management Console Support and downloads website (http://www14.software.ibm.com/webapp/set2/sas/f/hmcl/home.html).
2. Download the appropriate HMC V9 network images and save them on an FTP server.

   You cannot download these files directly to the HMC. You must download the image files to a server that accepts FTP requests.
3. Ensure that you download the following files:
   - img2a
   - img3a
   - base.img
   - disk1.img
   - hmcnetworkfiles.sum
4. Save the upgrade data on the HMC. Run the following commands to save the upgrade data:
   - To save data on both DVD and HDD, run the following commands:

     ```
     mount /media/cdrom
     ```

     ```
     saveupgdata -r diskdvd
     ```
   - To save data on the HDD, run the following command:

     ```
     saveupgdata -r disk
     ```
5. Copy the upgrade files to the bootable disk partition on the HMC. Run the **getupgfiles** command to copy the files.

   Example: **getupgfiles -h <ftp server> -u <user id> -d <remote directory>**

   Where,
   - **ftp server** is the host name or IP address of the FTP server where you download the HMC network images.
   - **user id** is a valid user ID on the FTP server. If you do not specify the password with the `--passwd` argument, you are prompted for a password.
   - **remote directory** is the directory on your FTP server where the HMC network images are saved.
6. Restart the HMC to upgrade the code that is copied to the bootable disk partition. Run the **chhmc -c altdiskboot -s enable --mode upgrade** to restart the HMC.
7. Restart the HMC and start the upgrade. Run the **hmcshutdown -r -t now** command to start the upgrade.

Securing the HMC

Learn how to enhance the security of your Hardware Management Console (HMC) that is based on your corporate security standards.

The default configuration of the HMC provides ample security for most enterprise users. With the Hardware Management Console (HMC) Version 8.4.0, or later, you can further enhance the security of the HMC that is based on your corporate security standards. To enhance the security for the HMC, you must set the HMC to a minimum of Level 1 security. You may choose Level 2 and Level 3 security depending on your environment and the corporate security requirements.

Note: Before changing the security level, check with your corporate security compliance team.

Level 1 security

To secure the HMC (level 1 security), complete the following steps:

1. Change the predefined password for the default `hscroot` user. For more information about password policy, see "Enhanced password policy" on page 86.
2. If the HMC does not belong to a physically secure environment, set the `grub` password by running the following command: `chhmc -c grubpasswd -s enable --passwd <new grub password>`
3. If you have configured the Integrated Management Module (IMM) on the HMC, set a strong IMM password.
4. Set a strong password for the *admin* user and general users on all servers.
5. Update the HMC with the latest released security fixes. For more information about the security fixes, see IBM Fix Central.

Level 2 security

If you have multiple users, complete the following steps to enhance the security for the HMC:

1. Create an account for each user on the HMC and assign the required roles and resources to users. For more information about the various roles in the HMC, see HMC tasks, user roles, IDs, and associated commands.

   Note: Ensure that you assign only the required resources and roles for users that are created on HMC. If necessary, you can also create custom roles.
2. Enable user data replication between different Hardware Management Consoles. The user data replication can be performed in Primary HMC-Secondary HMC mode or Peer-Peer mode. For more information about user data replication, see Manage Data Replication.
3. Import a certificate that is signed by the Certificate Authority.

Level 3 security

If you have multiple Hardware Management Consoles and system administrators, complete the following steps to enhance the security for the HMC:

1. Use centralized authentication such as Lightweight Directory Access Protocol (LDAP) or Kerberos. For more information about configuring LDAP, see How to Configure LDAP on HMC.
2. Enable user data replication between different Hardware Management Consoles.
3. Ensure that the HMC is in NIST SP 800-131A mode so that the HMC uses only strong ciphers.
4. Block ports that are not required in the firewall. For information about the HMC ports that can be used, see the following table:

Table 32. Port used by the user for interaction with HMC

Port	Description	Type	Protocol version (Default mode)	Protocol Version (NIST mode)
22	Open SSH	TCP	SSH v3	SSH v3
123	NTP	UDP	NTP	NTP
161	SNMP Agent	UDP	SNMP v3	SNMP v3
162	SNMP Trap	UDP	SNMP v3	SNMP v3
427	SLP	UDP	N/A	N/A

Table 32. Port used by the user for interaction with HMC (continued)

Port	Description	Type	Protocol version (Default mode)	Protocol Version (NIST mode)
443	HMC GUI and REST API	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
657	RMC	TCP/UDP	RSCT (Plain text + hash and sign)	RSCT (Plain text + hash and sign)
2300	5250 Terminal for IBM i	TCP	Plain text	Plain text
2301	5250 Secure terminal for IBM i	TCP	TLS 1.2	TLS 1.2
5989	CIM (legacy port, non-functional)	TCP	Non-functional	Non-functional
9900	FCS: HMC-HMC discovery	UDP	FCS	FCS
9920	FCS: HMC-HMC communication	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
9960	VTerm applet in GUI	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
12443	HMC REST API (legacy port)	TCP	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
12347	RSCT Peer Domain	UDP	RSCT (Plain text + hash and sign)	RSCT (Plain text + hash and sign)
12348	RSCT Peer Domain	UDP	RSCT (Plain text + hash and sign)	RSCT (Plain text + hash and sign)

Notes:

- You must use only SSH (port 22), HTTPS (port 443 and port 12443), 5250 secure terminal for IBM i (port 2301), and VTerm (port 9960) that are exposed to an intranet. All other ports must be used in a private or isolated network. You can use a separate Ethernet port and VLAN for the Resource Monitoring and Control (RMC) (port 657), FCS (port 9900 and port 9920), and RSCT Peer Domain (port 12347 and port 12348).
- Ports listed in the **`netstat`** command are used for internal processes only.

Enhanced password policy

You can enforce password requirements for locally authenticated users by using the Hardware Management Console (HMC). The enhanced password policy function allows the system administrator to set password restrictions. The enhanced password policy applies to the systems in which an HMC is installed.

System administrators can use the enhanced password policy to define a single password policy for all users. The HMC provides a medium security password policy, which can be activated by the system administrators to set password restrictions. The system administrator can also choose to activate the medium security policy or a new user-defined policy. The HMC medium security password policy cannot be removed from the system. The following table lists the attributes of the medium security policy and the default values.

Table 33. Password attributes for the HMC medium security password policy

Attribute	Description	Default value
`min_pwage`	The minimum number of days for which a password must remain active.	1
`pwage`	The maximum number of days for which a password might remain active.	365
`min_length`	The minimum length of a password.	8
`hist_size`	The number of previously saved passwords that cannot be reused.	10
`warn_pwage`	When the password is about to expire, the number of days before which a user is warned that the password is about to expire.	7
`min_digits`	The minimum number of digits that are required to be used in the password.	None
`min_uppercase`	The minimum number of upper case characters.	1
`min_lowercase`	The minimum number of lower case characters.	6
`min_special_chars`	The minimum number of special characters that must be used in the password.	None
`inactivity_expiration`	The number of days that can elapse before an HMC user account is disabled due to inactivity.	180

Consider the following items about the HMC medium security password policy:

- The policy features for password age, login disablement and inactivity expiration are not applicable to the **hscroot**, **hscpe**, and **root** user IDs. The password character validation is applicable for these user IDs.
- The policy affects only the locally authenticated users that are managed by the HMC and the policy cannot be enforced on LDAP or Kerberos users.
- The HMC medium security password policy or the user-defined policy allows the system administrators to set password reuse restrictions.
- The HMC medium security password is read-only and the attributes of HMC medium security password cannot be changed. You can create a new user-defined password to set password restriction.

You can use the following commands to configure the HMC medium security password policy:

mkpwdpolicy
Imports the password policy from a file, which contains all the parameters, or creates a password policy.

lspwdpolicy
Lists all the available password policy profiles and searches for specific parameters. You can also view the password policy that is currently active.

rmpwdpolicy
Removes an existing inactive password policy.

Note: You cannot remove an active medium security policy and the default read-only password policy.

chpwdpolicy
Changes parameters of an inactive password policy.

Solving common problems while securing the HMC

Learn how to solve some problems that you might encounter when you secure the HMC.

How to secure the connection between the Hardware Management Console (HMC) and the system?

The HMC connects to the system through the Flexible Service Processor (FSP). A proprietary binary protocol called Network Client protocol (NETC) is used for managing both FSP and Power hypervisor. The following table lists ports that are used by the HMC:

Table 34. Ports on FSP that are used to interact with the HMC

Port on FSP	Description	Protocol version (Default mode)	Protocol Version (NIST mode)
443	Advanced System Management Interface	HTTPS (TLS 1.2)	HTTPS (TLS 1.2)
30000	NETC	NETC (TLS 1.2). Falls back to SSLv3 for support of older firmware.	NETC (TLS 1.2)
30001	VTerm	NETC (TLS 1.2). Falls back to SSLv3 for support of older firmware.	NETC (TLS 1.2)

How to lock the HMC?

If you want to enhance the security for your infrastructure, you can use an Intrusion Prevention System (IPS) device or add all Hardware Management Consoles and IBM Power systems servers behind a firewall. Also, you can disable network services on the HMC if you do not use it remotely or if you want to lock the HMC down. To disable network services on the HMC, complete the following steps:

1. Disable remote command execution by using the SSH port.
2. Disable remote virtual terminal (VTerm port).
3. Disable remote web access (HMC graphical user interface and REST API).
4. Block ports in firewall by using HMC network settings for each configured Ethernet port.

How to set the HMC in NIST SP 800-131A compliance mode?

With HMC Version 8.1.0, or later, when you set the HMC in the compliance mode, only strong ciphers listed by NIST SP 800-131A are supported. You might not be able to connect to older Power systems servers such as, POWER5 servers that do not support Transport Layer Security (TLS 1.2). For more information about changing the security mode, see HMC V8R8 NIST mode.

How to view and change ciphers that are used by the HMC?

With HMC Version 8.1.0, or later, the HMC supports more secure cipher sets that are defined in NIST 800-131A. Ciphers that are used in the default mode are strong. For more information about encryption ciphers that are used by the HMC, run the **lshmcencr** command. If your corporate standards requires the use of a different set of ciphers, run the **chhmcencr** command to modify the encryption ciphers.

To list the encryption ciphers that are used by the HMC to encrypt user password, run the following command:

```
lshmcencr -c passwd -t c
```

To list the encryption ciphers that can currently be used by the HMC web user interface and REST API, run the following command:

```
lshmcencr -c webui -t c
```

To list the encryption ciphers and MAC algorithm that can currently be used by the HMC SSH interface, run the following command:

```
lshmcencr -c ssh -t c
```

```
lshmcencr -c sshmac -t c
```

How to check the strength of the certificate on the HMC?

The self-signed certificates on the HMC use SHA256 with 2048-bit RSA encryption, which is strong. If you are using CA signed certificates, ensure that you are not using the 1024-bit encryption, which is weak. The following certificates can be used for the HMC:

- The CA signed certificate can be used for the HMC graphical user interface and REST API (ports 443 and 12443).
- The port 9920 is used for HMC to HMC communication. You cannot replace this certificate with your own certificate.

How to choose between a self-signed certificate (default) or a CA signed certificate?

The HMC auto-generates a certificate during installation. However, you can generate a Certificate Signing Request (CSR) from the HMC and get a new certificate that is issued by a Certificate Authority. You can import this certificate into HMC. Ensure that you also obtain a domain name for the HMC. For more details about managing the certificates in HMC, see Manage Certificates.

How to audit the HMC?

The audit on the Hardware Management Consoles focuses on configured ciphers and the usage activity of the various HMC users. Use the following commands to view the usage activity of various HMC users:

Table 35. Ciphers that are used by the HMC

Purpose	Command
Password encryption (global setting)	`lshmcencr -c passwd -t c`
Password encryption for each user	`lshmcusr -Fname:password_encryption`
SSH ciphers	`lshmcencr -c ssh -t c`
SSH MAC	`lshmcencr -c sshmac -t c`
Cipher that are used for the HMC graphical user interface and REST API	`lshmcencr -c webui -t c`

Use the following commands to monitor various console and serviceable events information for uses in the HMC:

Table 36. Commands to view the logged on users and console or serviceable events information in the HMC

Information	Command
GUI users	`lslogon –r webui –u`
GUI tasks	`lslogon –r webui –t`

Table 36. Commands to view the logged on users and console or serviceable events information in the HMC (continued)

Information	Command
CLI users	`lslogon –r ssh –u`
CLI tasks	`lslogon –r ssh –t`
Operations on HMC	`lssvcevents -t console -d <number of days>`
Operations on System	`lssvcevents –t hardware –m <managed system> -d <number of days>`

Centralized monitoring events for the HMC: If you have many Hardware Management Consoles, set the `rsyslog` file to collect all the usage data.

How does IBM fix the HMC security vulnerabilities?

IBM has a security incidence response process named IBM Product Security Incident Response Team (PSIRT). The IBM Product Security Incident Response Team (PSIRT) is a global team that manages the receipt, investigation and internal coordination of security vulnerability information related to IBM offerings. Open Source and IBM components that are shipped with the HMC are actively monitored and analyzed. Interim fixes and security fixes are provided by IBM for all supported releases of the HMC.

How to track new interim fixes on the HMC?

The security bulletin contains information about the vulnerability and interim fixes for supported HMC versions. To track interim fixes on the HMC, you can:

- Search the latest security bulletins at IBM Security Bulletin.
- Follow @IBMPowereSupp on Twitter for notifications.
- Subscribe to email notifications at IBM Support.

Security profiles: Global Data Protection Regulation (GDPR) and Payment Card Industry Data Security Standard (PCI-DSS)

Learn about how the Hardware Management Console (HMC) handles the privacy information of the users.

The Hardware Management Console (HMC) is a closed appliance that does not store any cardholder data. Hence, only a subset of requirements and security assessment procedures of IT security that are defined by PCI-DSS are applicable for the HMC. Only trusted code that is distributed by IBM can be installed on the HMC. When any vulnerability is known through the IBM PSIRT process, interim fixes are published. The requirements and recommendations include the following items:

GDPR queries

Table 37. GDPR queries . The table provides information about the questions related to GDPR.

Questions	Answers
What kind of data is stored in the HMC?	HMC stores configuration information of Power hardware, PowerVM virtualization, and the performance metrics information.
Does the HMC process any personal data?	You can provide contact information for the call home function. Providing contact information for the call home function is optional.

Table 37. GDPR queries . The table provides information about the questions related to GDPR. *(continued)*

Questions	Answers
Which predefined accounts are used for system administration of the HMC?	The system administrator user uses the *hscroot* username.
Do any of the accounts in the HMC relate to a specific person?	No.
Is it mandatory to provide personal data in the HMC?	No. You do not need to provide personal data information. However, providing this information is optional.
Does the HMC log file have any personal data information?	No.
Is it possible to delete personal data completely and permanently?	Yes. Unconfigure the call home function.

PCI-DSS queries

Table 38. PCI-DSS queries . The table provides information about the questions related to PCI-DSS

Questions	Answers
How to install and maintain a firewall configuration to protect the data of the cardholder?	The HMC does not store or access any cardholder data. However, the HMC has a firewall configuration and the user can control and enable specific ports.
Can I use vendor-supplied default value for system passwords and other security parameters?	Before you install a system on the network, change all the predefined passwords of the *hscroot* user.
How does the HMC protect the stored data of the cardholder?	The HMC does not store or access any cardholder data.
How does the HMC encrypt the data of the cardholder when the data is transmitted across open public networks?	The HMC does not store or access any cardholder data.
How to use and regularly update anti-virus software programs?	The HMC is a closed appliance. Therefore, malware cannot infect the HMC.
How to develop and maintain secure systems and applications?	You must install the required patches to your system manually from the IBM Fix Central website. Only trusted code that are distributed by IBM can be installed on the HMC.
Does the HMC restrict access to the cardholder data?	The HMC does not store or access any cardholder data.
How to assign a unique ID to each person who has access to the computer?	You can implement this requirement by ensuring that there are no shared IDs and by following the password policies.
How to restrict the physical access to the data of the cardholder?	The HMC does not store or access any cardholder data
How to track and monitor the access to network resources and to the cardholder data?	The HMC does not store or access any cardholder data.

Table 38. PCI-DSS queries . The table provides information about the questions related to PCI-DSS *(continued)*	
Questions	**Answers**
How does the HMC test the security of the system and processes?	Scan tools are used to run security scans on all the released versions of the HMC. The scan tools include: *Qualys*, *Nessus*, *testssl*, *sslscan* and *ASoC*.
How to maintain a security policy that includes information security for employees and contractors?	System administrator disables the remote user login, activates the user login on a need basis, and deactivates the user login when the access is no longer required.

HMC port locations

You can find port locations by using location codes. Use the HMC port location illustrations to map a location code to the HMC port position on the server.

Model 9080-HEX HMC port locations

Use this diagram and table to map the HMC ports on the 9080-HEX.

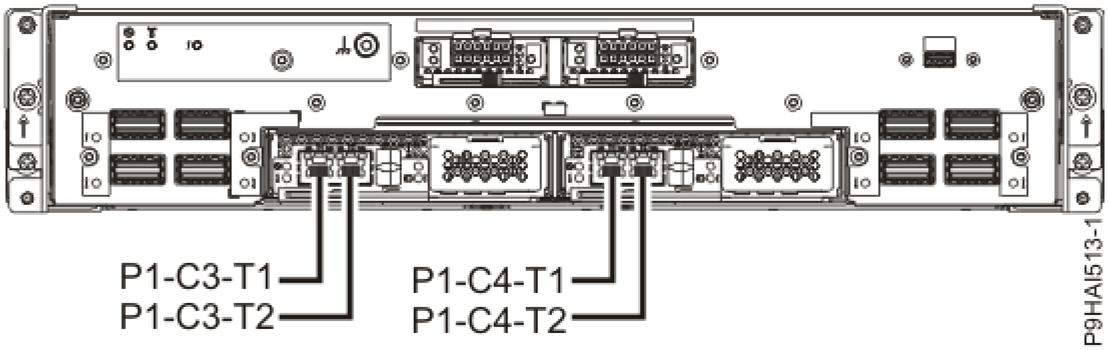

Figure 6. 9080-HEX HMC port locations

Table 39. 9080-HEX HMC port locations		
Port	**Physical port location**	**Identify LED**
Service processor card 1 - HMC port 1	Un-P1-C3-T1	No
Service processor card 1 - HMC port 2	Un-P1-C3-T2	No
Service processor card 2 - HMC port 1	Un-P1-C4-T1	No
Service processor card 2 - HMC port 2	Un-P1-C4-T2	No
For more information about HMC port locations on the 9080-HEX , see Part location and location codes.		

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with ICT Accessibility 508 Standards and 255 Guidelines (https://www.access-board.gov/ict/) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power servers.

The IBM Power servers online product documentation in IBM Documentation is enabled for accessibility. For more information about IBM's commitment to accessibility, see the IBM accessibility website at IBM Accessibility (https://www.ibm.com/able/).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name and IP address for purposes of session management. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Electronic emission notices

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Class A Notices

The following Class A statements apply to the IBM servers that contain the Power10 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

The following Class A statements apply to the servers.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

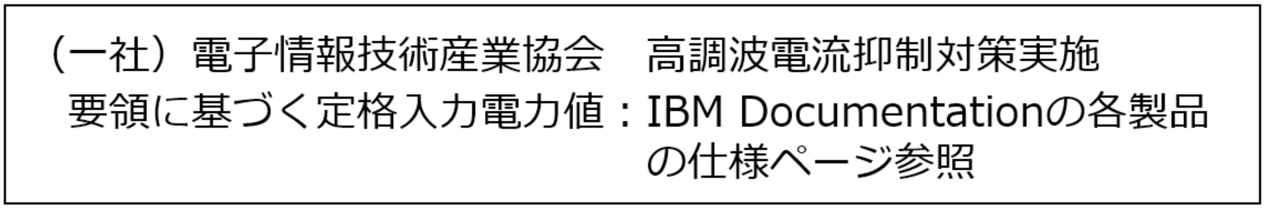

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：IBM Documentationの各製品
の仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

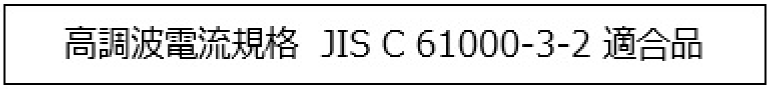

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

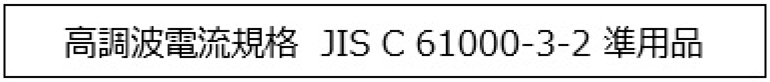

高調波電流規格 JIS C 61000-3-2 準用品

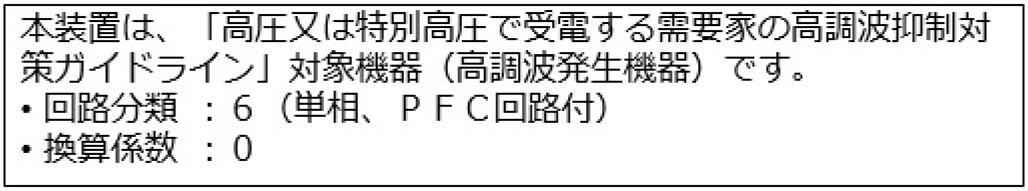

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

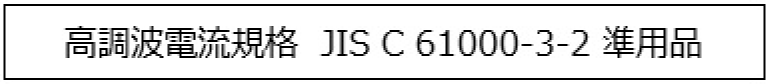

高調波電流規格 JIS C 61000-3-2 準用品

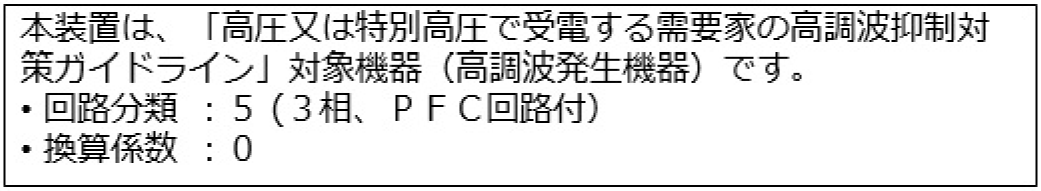

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

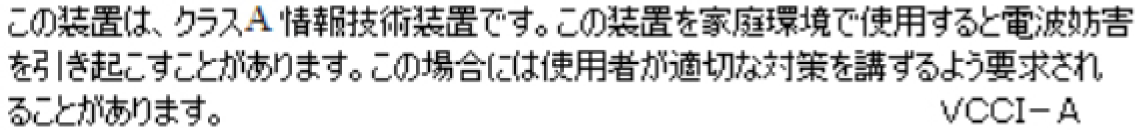

この装置は、クラスA 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI－A

Korea Notice

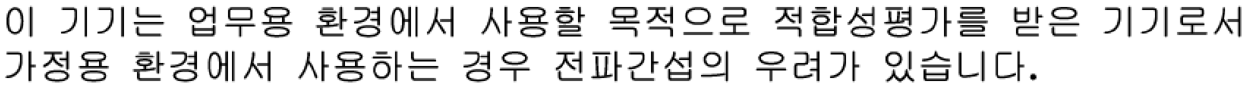

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

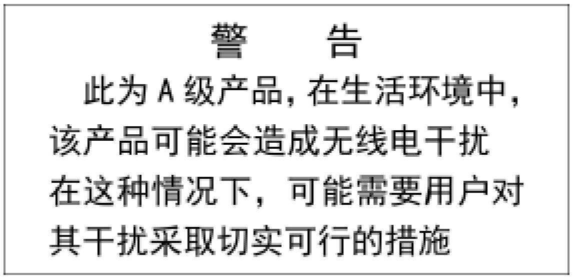

警　告

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰
在这种情况下，可能需要用户对
其干扰采取切实可行的措施

Russia Notice

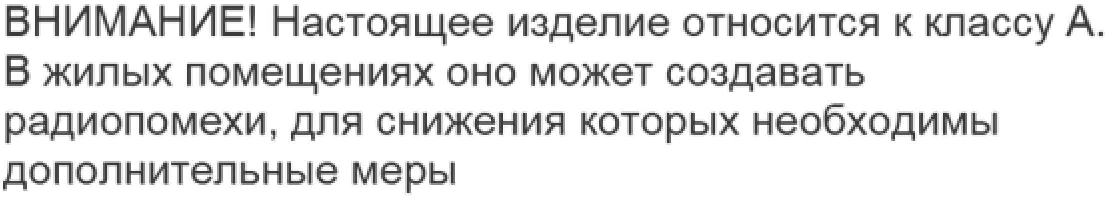

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

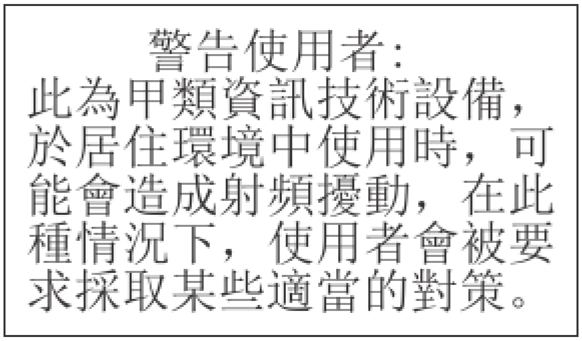

警告使用者:
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

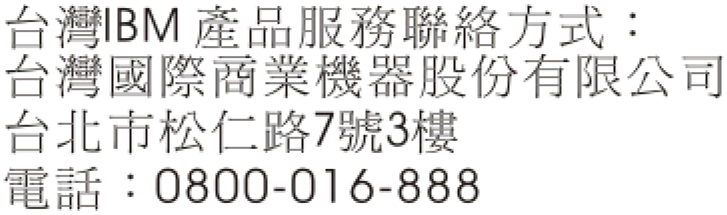

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

United Kingdom Notice

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

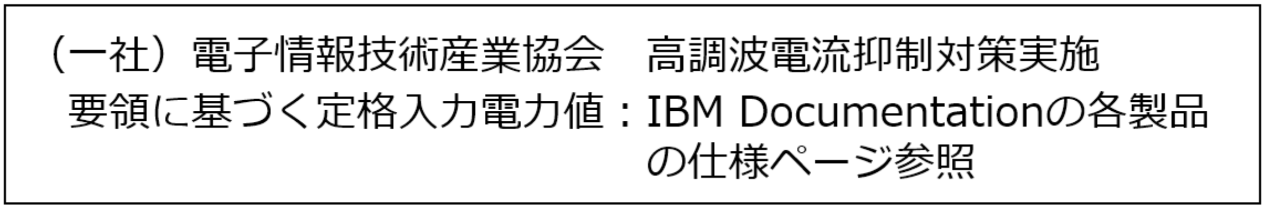

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：IBM Documentationの各製品
の仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

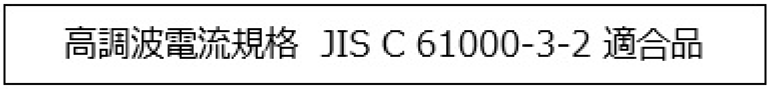

高調波電流規格　JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

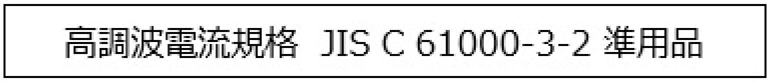

高調波電流規格　JIS C 61000-3-2 準用品

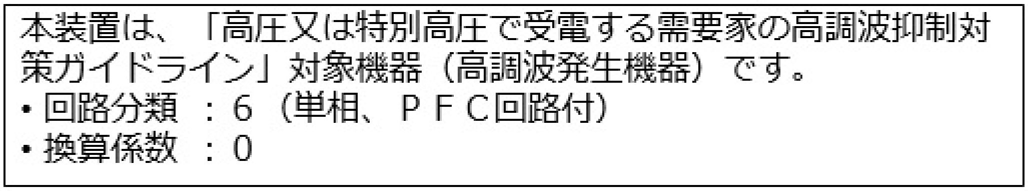

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
• 回路分類 ：6（単相、ＰＦＣ回路付）
• 換算係数 ：0

This statement applies to products greater than 20 A per phase, three-phase.

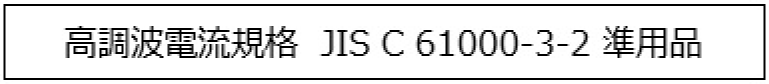

高調波電流規格 JIS C 61000-3-2 準用品

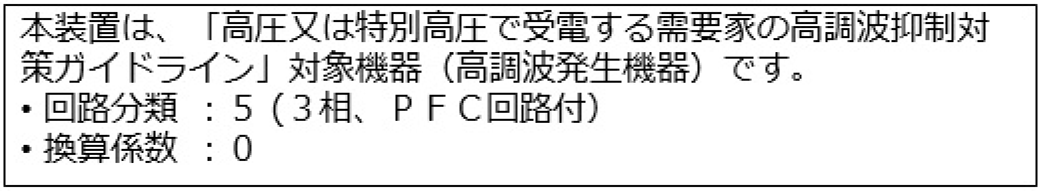

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

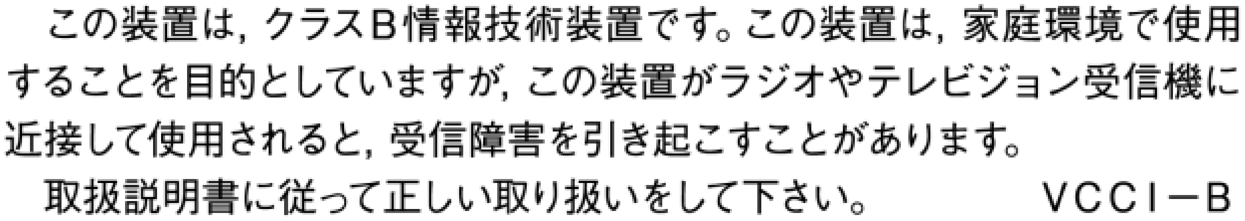

この装置は，クラスＢ情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。 ＶＣＣＩ－Ｂ

Taiwan Notice

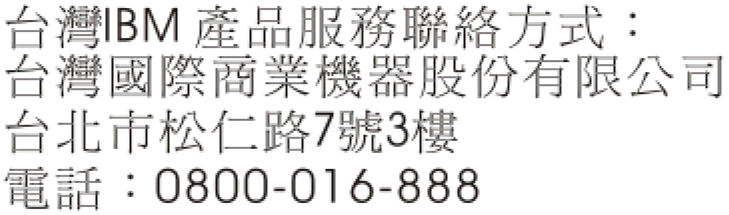

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Managing the Hardware Management Console

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 109.

This edition applies to IBM Hardware Management Console Version 10 Release 1 Maintenance Level 1010 and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Managing the HMC

Learn how to use the Hardware Management Console (HMC).

About this task

Learn about the tasks that you can use on the console and how to navigate by using the web-based user interface with graphical views of managed systems and simplified navigation.

Note: Many of the HMC tasks that are listed here can also be performed by using PowerVC. For more information about the tasks that you can perform by using PowerVC, see HMC and PowerVC.

What's new in Managing the HMC

Read about new or significantly changed information in Managing the HMC since the previous update of this topic collection.

September 2021

- Updated the menu icons for the HMC Graphical User Interface.
- Updated the following topics:
  - “Using the web-based user interface” on page 3
  - “Overview of menu options” on page 4
  - “Power Management” on page 31
  - “Processor, Memory, I/O” on page 42
  - “PowerVM” on page 44
  - “Partition Properties” on page 56
  - “Schedule Operations” on page 59
  - “Validate Maintenance Readiness and Prepare” on page 60
  - “Console Default Settings” on page 71
  - “Manage Data Replication” on page 75
  - “Manage Virtual I/O Server Backups” on page 78
  - “Manage KDC” on page 86

Introduction to the HMC

Learn about some of the concepts and functions of the Hardware Management Console (HMC) and the user interface that is used for accessing those functions.

You can configure and manage servers on the HMC. One HMC can manage multiple servers, and dual HMCs can provide redundant support by managing the same system. To ensure consistent function, each HMC is shipped preinstalled with the HMC Licensed Machine Code Version 9, Release 1.

To provide flexibility and availability, you can implement HMCs in several configurations.

HMC as the DHCP server

An HMC that is connected by either a private network to the systems it manages might be a DHCP server for the service processors of the systems. An HMC might also manage a system over an open network, where the managed system's service processor IP address is assigned by a customer-supplied DHCP server or manually assigned by using the Advanced System Management Interface (ASMI).

Physical proximity

Before HMC Version 7, at least one local HMC was required to be physically located near the managed systems. As an alternative to the local HMC, you can use a supported device, such as a personal computer that has connectivity and authority to operate through a remotely attached HMC. The local device must be in the same room as your server and at a distance of 8 m (26 ft) from your server. The local device must have the functional capability that is equivalent to the HMC that it replaces and that is needed by the service representative to service the system. For a virtual HMC, the functional capabilities also include the method of transferring service data, such as firmware updates or diagnostic data, and transferring the log information to and from the HMC.

Redundant or Dual HMCs

A server might be managed by either 1 or 2 Hardware Management Consoles. When two Hardware Management Consoles manage one system, they are peers, and each HMC can be used to control the managed system. The best practice is to attach one HMC to the supported networks or HMC ports of the managed systems. The networks are intended to be independent. Each HMC might be the DCHP server for a service network. Because the networks are independent, the DHCP servers must be set up to provide IP addresses on two unique and nonroutable IP ranges.

Redundant or Dual HMCs that manage the same server must not be at different version and release levels. For example, an HMC at Version 7 Release 7.1.0 and an HMC at Version 7 Release 3.5.0 cannot manage the same server. The HMCs must be at the same version and release level.

When the server is connected to the higher version of the management console, the partition configuration is upgraded to the latest version. After the partition configuration upgrade, lower levels of the management consoles will not be able to interpret the data correctly. After the server is managed by the higher version of the management console, you must first initialize the server before you can go back to the lower version of the management console. You can restore a backup that is taken at the older level or re-create the partitions. If the server is not initialized, one of the following outcomes can occur depending on the version of the lower-level HMC:

- HMC Version 7 Release 7.8.0 and later reports a connection error of **Version mismatch** with reference code **Save Area Version Mismatch**.
- HMC Version 7 Release 7.7.0 and earlier might report a server state of **Incomplete** or **Recovery**. In addition, partition configuration corruption can also occur.

Logging on to the Hardware Management Console (HMC) when PowerSC MFA is configured on the HMC

Learn how to log in to the HMC when IBM PowerSC Multi-factor Authentication (MFA) is configured on the HMC.

If IBM PowerSCMFA is enabled on the HMC and the user is configured on the PowerSC MFA server, you can choose to log in to the HMC by first entering the user ID and a policy name that is provided by your security administrator. You are then prompted to provide additional credentials.

In the HMC login page, if you click **Policy Name**, the authentication mechanism is set to the in-band authentication type. For example, if the policy that you want to use is associated with the Rivest-Shamir-Adleman (RSA) authentication method, you can enter the secure ID passcode that you received from the RSA secure ID device or the application. Then, click **Next or Sign In** to log in to the HMC.

Notes:

- If MFA is not enabled on the HMC, you can log in to the HMC with the user ID and password.
- If you obtain a cache token credential (CTC) code from the PowerSC MFA server that is configured by your security administrator, enter the CTC code in the **Password** field.

Predefined user IDs and passwords

Predefined user IDs and passwords are included with the Hardware Management Console (HMC). It is imperative to the security of your system that you change the `hscroot` predefined password immediately.

If the password expires when you try to log in to the HMC, complete the following steps:

1. Enter the **Current Password** and the **New Password**.
2. Re-enter the new password in the **Confirmation new password** field.
3. Click **OK**. If the new password complies with the current password policy, the password for the HMC is changed.

The following predefined user IDs and passwords are included with the HMC:

Table 1. Predefined HMC user IDs and passwords

User ID	Password	Purpose
`hscroot`	`abc123`	The `hscroot` user ID and password are used to log in to the HMC for the first time. They are case-sensitive and can be used only by a member of the super administrator role.

Using the web-based user interface

You can use the web-based user interface to perform tasks on the Hardware Management Console (HMC) or on your managed resources.

This user interface comprises several major components: the title bar, the navigation area, the content pane, the menu pod, and the dock pod.

The *title bar*, across the top of the workplace window, identifies the product, help options, task log, serviceable events, and any user that is logged in.

The *navigation area*, in the left portion of the window, contains the primary navigation links for selecting your system and starting tasks for your HMC.

The *content pane*, in the middle portion of the window, displays information that is based on the current selection from the navigation area. For example, when **All Systems** is selected by clicking on **Resources** in the navigation area, all the available systems are shown in the content pane.

The *menu pod*, in the left portion of the window, is displayed after you select a system and provides quick access to commonly used HMC tasks and views of resources and properties.

You can resize the panes of the HMC workplace by moving the mouse pointer over the border that separates the navigation pane from the work pane until the mouse pointer changes to a double-pointed arrow. When the pointer changes shape, press and hold the left mouse button while you drag the mouse pointer to the left or right. Release the button and your navigation pane or work pane is now larger or smaller in size. You can also do this task within the work pane border that separates the resources table from the taskpad.

You can change the layout of the *content pane* according to your preference by clicking **Display Gallery View**, **Display Table View**, or **Display Relationship View**.

You can reposition columns in tables by selecting and dragging a column to a new position. You can also select which columns to display by clicking the drop-down menu that is located next to the last column of each table. You can save your preferences by clicking the **Save User Preferences** icon.

You can change how many rows are displayed in tables on each page by clicking one of the **Items per page** (**10**, **20**, **30**, or **50**) icons that are located below each table.

Note: Pop-up windows must be enabled for full functionality of the HMC.

Overview of menu options

Learn about the menu options and associated tasks that are available in the Hardware Management Console (HMC).

The menu options and tasks that are described in this section are available in the HMC interface.

Table 2. HMC menu options

Menu	Submenu	Options/Tasks
Resources	All Systems	View All Systems
	All Partitions	View All Partitions
	All Virtual I/O Servers	View All Virtual I/O Servers
	All Power Enterprise Pools	View All Power Enterprise Pools
	All Shared Storage Pool Clusters	View All Shared Storage Pool Clusters
	All Groups	View All Groups

Table 2. HMC menu options (continued)

Menu	Submenu	Options/Tasks
Console Management	Console Settings	Launch Guided Setup Wizard
		View Network Topology
		Test Network Connectivity
		Change Network Settings
		Change Performance Management Settings
		Change Date and Time
		Change Language and Locale
	Console Management	Shut Down or Restart the Management Console
		Schedule Operations
		View Licences
		Update the Hardware Management Console
		Manage Install Resources
		Manage Virtual I/O Server Image Repository
		Format Media
		Backup Management Console Data
		Restore Management Console Data
		Save Upgrade Data
		Manage Data Replication
	Template Library	System and Partition Library
	Updates	Not available (use the Update the Hardware Management Console option instead)

Table 2. HMC menu options (continued)

Menu	Submenu	Options/Tasks
Users and Security	Users and Roles	Change User Password
		Manage User Profiles and Access
		Manage Users and Tasks
		Manage Task and Resource Roles
	Systems and Console Security	Manage Certificates
		Manage LDAP
		Manage KDC
		Enable Remote Command Execution
		Enable Remote Operation
		Enable Remote Virtual Terminal

Table 2. HMC menu options (continued)

Menu	Submenu	Options/Tasks
Serviceability	Console Events Logs	View Console Events window
	Serviceable Events Manager	Serviceable Events Manager window
	Events Manager for Call Home	Events Manager for Call Home window
	Service Management	Create Serviceable Event
		Manage Remote Connections
		Manage Remote Support Requests
		Manage Dumps
		Transmit Service Information
		Schedule Service Information
		Format Media
		Perform Management Console Trace
		View Management Console Logs
		View Component Logs
		Electronic Service Agent Setup Wizard
		Authorize User
		Enable Electronic Service Agent
		Manage Outbound Connectivity
		Manage Inbound Connectivity
		Manage Customer Information
		Manage Serviceable Event Notification
		Manage Connection Monitoring

Tasks and roles

Each HMC user can be a member of a different role. Each of these roles allows the user to access different parts of the HMC and complete different tasks on the managed system. HMC roles are either predefined or customized.

The roles that are discussed refer to HMC users; operating systems that are running on logical partitions have their own set of users and roles. When you create an HMC user, you must assign that user a task role. Each task role allows the user different levels of access to tasks available on the HMC interface. For more information about the tasks each HMC user role can perform, see "HMC tasks, user roles, IDs, and associated commands" on page 8.

You can assign managed systems and logical partitions to individual HMC users. This action allows you to create a user that has access to managed system A but not to managed system B. Each grouping of managed resource access is called a managed resource role.

The **predefined** HMC roles, which are the default on the HMC, are as follows:

Table 3. Predefined HMC Roles

Role	Description	HMC User ID
Operator	The operator is responsible for daily system operation.	**hmcoperator**
Super administrator	The super administrator acts as the root user, or manager, of the HMC system. The super administrator has unrestricted authority to access and modify most of the HMC system.	**hmcsuperadmin**
Product engineer	A product engineer helps support situations, but cannot access HMC user management functions. To provide support access for your system, you must create and administer user IDs with the product engineer role.	
Service representative	A service representative is an employee who is at your location to install, configure, or repair the system.	**hmcservicerep**
Viewer	A viewer can view HMC information, but cannot change any configuration information.	**hmcviewer**
Client live update	The client live update role is intended for use when you are using the AIX® Live Update capability on a partition of a managed system. A client live update user has authority that is limited to what is necessary to complete a live update on AIX.	**hmcclientliveupdate**

You can create **customized** HMC roles by modifying predefined HMC roles. Creating customized HMC roles is useful for restricting or granting specific task privileges to a certain user.

HMC tasks, user roles, IDs, and associated commands

The roles discussed in this section refer to HMC users; operating systems running on logical partitions has its own set of users and roles.

Each HMC user has an associated task role and a resource role. The task role defines the operations the user can perform. The resource role defines the systems and partitions for performing the tasks. The users may share task or resource roles. The HMC is installed with five predefined task roles. The single predefined resource role allows access to all resources. The operator can add customized task roles, customized resource roles, and customized user IDs.

Some tasks have an associated command. For more information about accessing the HMC command line, see “Using the HMC remote command line” on page 101.

Some tasks can only be performed using the command line. For a listing of those tasks, see Table 9 on page 25.

For more information about where to find task information, see the following table:

Table 4. HMC task groupings

HMC tasks and the corresponding user roles, IDs, and commands	Associated table
HMC Management	Table 5 on page 9
Service Management	Table 6 on page 12
Systems Management	Table 7 on page 14
Control Panel Functions	Table 8 on page 23

This table describes the HMC management tasks, commands, and default user roles associated with each HMC Management task.

Table 5. HMC Management tasks, commands, and default user roles

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Backup Management Console Data" on page 74 bkconsdata	X	X		X
Backup Profile Data bkprofdata	X	X		X
Change BMC Certificates chbmccert	X	X		X
Certificate Management chhmccert lshmccert mkhmccert		X		
"Change Date and Time" on page 70 chhmc lshmc	X	X		X
"Change Language and Locale" on page 71 chhmc lshmc	X	X	X	X
Change HMC Configuration chipsec chpsm chusrtca	X	X		X

Table 5. HMC Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Change Network Settings" on page 69 chhmc lshmc	X	X		X
Change Proxy Configuration chproxy		X		X
"Change User Password" on page 81 chhmcusr	X	X	X	X
List BMC Certificates lsbmccert	X	X	X	X
List HMC Configuration lsipsec lspsm lsusrtca	X	X	X	X
List HMC Encryption Task lshmcencr	X	X	X	
List System Plan lssysplan		X		
List Proxy Configuration lsproxy	X	X	X	X
"Manage KDC" on page 86 chhmc lshmc getfile rmfile		X		
"Manage LDAP" on page 86 lshmcldap chhmcldap		X		
"Launch Guided Setup Wizard" on page 67		X		

Table 5. HMC Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Launch Remote Hardware Management Console	X	X	X	X
Lock HMC Screen	X	X	X	X
Logoff or Disconnect	X	X	X	X
"Manage Certificates" on page 85		X		
"Manage Data Replication" on page 75	X	X		
"Manage Task and Resource Roles" on page 84 chaccfg lsaccfg mkaccfg rmaccfg		X		
"Manage User Profiles and Access" on page 81 chhmcusr lshmcusr mkhmcusr rmhmcusr		X		
"Manage Users and Tasks" on page 83 lslogon termtask	X	X	X	X
Open 5250 Console	X	X		X
"Enable Remote Command Execution" on page 90 chhmc lshmc	X	X		X
"Enable Remote Operation" on page 90 chhmc lshmc	X	X	X	X

Table 5. HMC Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Enable Remote Virtual Terminal" on page 91	X	X		X
"Restore Management Console Data" on page 75	X	X		X
"Save Upgrade Data" on page 75 saveupgdata	X	X		X
"Schedule Operations" on page 72	X	X		
"Shut Down or Restart" on page 72 hmcshutdown	X	X		X
"Serviceable Events Manager" on page 47 lssvcevents	X	X		X
"View Licenses" on page 73	X	X	X	X

This table describes the Service Management tasks, commands, and default user roles.

Table 6. Service Management tasks, commands, and default user roles

HMC Interface Tasks and Associated Commands	User roles and IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Create Serviceable Event" on page 48		X		X
"Serviceable Events Manager" on page 92 chsvcevent cpfile lssvcevents mksvcevent updpmh		X		X

Table 6. Service Management tasks, commands, and default user roles (continued)				
HMC Interface Tasks and Associated Commands	**User roles and IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Format Media" on page 74 formatmedia	X	X		X
"Manage Dumps" on page 93 dump cpdump getdump lsdump startdump lsfru	X	X		X
"Transmit Service Information" on page 93 chsacfg lssacfg	X	X		
"Enable Electronic Service Agent" on page 95	X	X		X
"Manage Outbound Connectivity" on page 95	X	X		X
"Manage Inbound Connectivity" on page 96	X	X		X
"Manage Customer Information" on page 96	X	X		X
"Authorize User" on page 94		X		
"Manage Event Notification" on page 97 chsacfg lssacfg	X	X		X
"Manage Connection Monitoring" on page 97	X	X	X	X
"Electronic Service Agent Setup Wizard" on page 94		X		X

This table describes the Systems Management tasks, commands, and default user roles.

Table 7. Systems Management tasks, commands, and default user roles

HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
“General Settings” on page 41 lshwres	X	X	X	X
lsled	X	X	X	X
lslparmigr	X	X	X	X
lssyscfg	X	X	X	X
chhwres	X	X	X	X
chsyscfg	X	X	X	X
migrlpar	X	X	X	X
optmem	X	X		X
lsmemopt	X	X	X	X
lsrrstartlpar	X	X		
Update Password chsyspwd		X		
Change Default User Interface Settings	X	X	X	X
List CEC Property lscomgmt lsiotopo	X	X	X	X
List Utilization Data lslparutil	X	X	X	X
Operations				
“Power Off” on page 30 chsysstate	X	X		X
“Activate” on page 57 chsysstate	X	X		X
“Save Current Configuration” on page 63 chsysstate	X	X		X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Restart" on page 58 chsysstate	X	X		X
"Shut Down" on page 58 chsysstate	X	X		X
chlparstate	X	X		X
LED Status: Deactivate Attention LED "Attention LED" on page 34 chled	X	X		
LED Status: Identify LED "Attention LED" on page 34	X	X	X	X
LED Status: Test LED "Attention LED" on page 34	X	X	X	X
"Schedule Operations" on page 32	X	X		
"Launch ASM Interface" on page 33 asmmenu	X	X		X
"Rebuild" on page 33 chsysstate	X	X		
"Power Management" on page 31 chpwrmgmt lspwrmgmt		X		
"Delete" on page 59 rmsyscfg	X	X		X
"Mobility" on page 61 lslparmigr migrlpar	X	X		X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
"Manage Profiles" on page 62 chsyscfg lssyscfg mksyscfg rmsyscfg chsysstate	X	X		X
Manage System Plan cpsysplan rmsysplan		X		
Make System Plan mksysplan		X		
Deploy System Plan deploysysplan		X		
Change N_Port Login chnportlogin	X	X		X
RR Start LPAR lsrrstartlpar rrstartlpar	X	X		
Migrate LPAR migrdbg refdev	X	X		
Make Profile Data mkprofdata	X	X		
Restore Profile Data migrcfg	X	X		
Remove Profile Data rmprofdata	X	X		

Table 7. Systems Management tasks, commands, and default user roles (continued)				
HMC Interface Tasks and Associated Commands	**User roles/IDs**			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmi n)**	**Viewer (hmcview er)**	**Service Representative (hmcservicerep)**
Manage Pmem CEC Config: Initialize Profile Data: Restore Profile Data rstprofdata For option "--retainpmemvolume" (access only for hmcsuperadmin)	X	X		
Vios Admin Op: Virtual IO Server Command viosvrcmd For option "--admin" (access only for hmcsuperadmin and hmcoperator)	X	X		X
"Operations" on page 30	X	X	X	X
Configuration				
"Create Partition from Template" on page 36		X		
"Deploy System from Template" on page 36		X		
"Capture Configuration as Template" on page 36		X		
Change CEC Property chprimhmc	X	X		
Change Trusted System Key chtskey		X		
"Create Partition" on page 40		X		
List LPAR Property lsmigrdbg	X	X	X	X
Hibernate LPAR lsrsdevsize	X	X		
List N_Port Login lsnportlogin	X	X		X
LS Profile Space lsprofspace	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
List Trusted System Key lstskey	X	X	X	X
"Manage Custom Groups" on page 62	X	X		X
"Manage Profiles" on page 62 chsyscfg chsysstate lssyscfg mksyscfg rmsyscfg	X	X	X	X
Manage License Keys chlickey	X	X		
Manage Utilization Data chlparutil	X	X		X
Save Current Configuration "Save Current Configuration" on page 63 mksyscfg	X	X		
ViewSPP lsmemdev	X	X	X	X
Connections				
"Service Processor Status" on page 35 lssysconn	X	X	X	X
"Reset or Remove Connections" on page 35 rmsysconn	X	X		
Add Connection mksysconn	X	X		
Open V Term mkvterm	X	X		X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Close V Term rmvterm	X	X		X
“Disconnect Another HMC” on page 35		X		
Hardware (Information)				
“Hardware Operations” on page 49	X	X	X	X
Updates				
“Change Licensed Internal Code” on page 37 lslic updlic		X		X
“Check System Readiness” on page 37 updlic		X		X
“View System Information” on page 36 lslic		X		X
Update HMC updhmc lshmc		X		X
Serviceability				
“Serviceable Events Manager” on page 63 chsvcevent lssvcevents		X		X
Change SNMP Alerts chspsnmp	X	X		X
“Create Serviceable Event” on page 48		X		X
“Reference Code Log” on page 64 lsrefcode	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
"Control Panel Functions" on page 64 lssyscfg	X	X		
"Add FRU" on page 50		X		X
"Add Enclosure" on page 51		X		X
"Exchange FRU" on page 50		X		X
"Remove FRU" on page 51		X		X
"Remove Enclosure" on page 51		X		X
"Power On/Off Unit" on page 50		X		X
"Manage Dumps" on page 48 dump cpdump getdump lsdump startdump lsfru	X	X		X
"Collect VPD" on page 49	X	X	X	X
"Type, Model, Feature" on page 49		X		
"Setup FSP Failover" on page 52 chsyscfg lssyscfg		X		
"Initiate FSP Failover" on page 52 chsysstate		X		
List CEC Property lsprimhmc	X	X	X	X
Capacity on Demand (CoD)				
Enter CoD code chcod		X		

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
View History Log lscod	X	X	X	X
Change CEC Property chcomgmt	X	X		
CoD Pool Management: Change CoD chcodpool	X	X		
Change CoD mkcodpool		X		
Change VET Code chvet		X		
List CoD Information lscodpool	X	X	X	X
List VET Information lsvet	X	X	X	X
Processor: View Capacity Settings lscod	X	X	X	X
Processor CUoD: View Code Information lscod	X	X	X	X
Processor: On/Off CoD: Manage chcod		X		
Processor: On/Off CoD: View Capacity Settings lscod	X	X	X	X
Processor: On/Off CoD: View Billing Information lscod	X	X	X	X
Processor: On/Off CoD: View Code Information lscod	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
Processor: Trial CoD: Stop chcod		X		
Processor: Trial CoD: View Capacity Settings lscod	X	X	X	X
Processor: Trial CoD: View Code Information lscod	X	X	X	X
Processor: Reserve CoD: Manage chcod		X		
Processor: Reserve CoD: View Capacity Settings lscod	X	X	X	X
Processor: Reserve CoD: View Code Information lscod	X	X	X	X
Processor: Reserve CoD: View Shared Processor Utilization lscod	X		X	X
PowerVM (formerly known as Advanced POWER® Virtualization): Enter Activation Code chcod		X		
PowerVM: View History Log lscod	X	X	X	X
PowerVM: View Code Information lscod	X	X	X	X
Enterprise Enablement: Enter Activation Code chcod		X		
Enterprise Enablement: View History Log lscod	X	X	X	X

Table 7. Systems Management tasks, commands, and default user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Enterprise Enablement: View Code Information lscod	X	X	X	X
Other Advanced Functions: Enter Activation Code chcod		X		
Other Advanced Functions: View History Log lscod	X	X	X	X
Other Advanced Functions: View Code Information lscod	X	X	X	X
Processor: Manage chcod		X		
Processor: View Capacity Settings lscod	X	X	X	X
Processor: View Code Information lscod	X	X	X	X
Memory: Manage chcod		X		
Memory: View Capacity Settings lscod	X	X	X	X
Memory: View Code Information lscod	X	X	X	X

This table describes the Control Panel Functions tasks, commands, and default user roles.

Table 8. Control Panel Functions tasks, commands, and user roles

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	Super Administrator (hmcsuperadmin)	Viewer (hmcviewer)	Service Representative (hmcservicerep)
Serviceability				

Table 8. Control Panel Functions tasks, commands, and user roles (continued)

HMC Interface Tasks and Associated Commands	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadm in)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
(21) Activate Dedicated Service Tools chsysstate	X	X		
(65) Disable Remote Service chsysstate	X	X		
(66) Enable Remote Service chsysstate	X	X		
(67) DIsk Unit IOP Reset / Reload chsysstate	X	X		
(68) Concurrent Maintenance Power Off Domain	X	X		
(69) Concurrent Maintenance Power On Domain	X	X		
(70) IOP Control Storage Dump chsysstate	X	X		
(71) Product Engineering Debug Tools pedbg				
(72) PE Shell Access pesh	X	X	X	X

This table describes the commands that are not associated with an HMC UI task, and defines the default user roles that can perform each command.

Table 9. Command line tasks, associated commands, and user roles

Command line tasks	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
Change which encryption is used by the HMC to encrypt the passwords of locally authenticated HMC users, or change which encryptions can be used by the HMC Web UI. chhmcencr		X		
List which encryption is used by the HMC to encrypt the passwords of locally authenticated HMC users, or list which encryptions can be used by the HMC Web UI chhmcfs	X	X	X	
Free up space in HMC file systems chhmcfs	X	X		
List HMC file system information lshmcfs	X	X	X	X
Test for removable media readiness on the HMC ckmedia	X	X		X
Obtain required files for an HMC upgrade from a remote site getupgfiles	X	X		X
Provide screen capture on the HMC hmcwin	X	X	X	X
Log SSH command usage logssh	X	X	X	X
Clear or dump partition configuration data on a managed system lpcfgop		X		

Table 9. Command line tasks, associated commands, and user roles (continued)

Command line tasks	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
List environmental information for a managed frame, or for systems contained in a managed frame lshwinfo	X	X	X	X
List which HMC owns the lock on a managed frame lslock	X	X	X	X
Force an HMC lock on a managed frame to be released rmlock		X		
List the storage media devices that are available for use on the HMC lsmediadev	X	X	X	X
Manage SSH authentication keys mkauthkeys	X	X	X	X
Monitoring HMC subsystems and system resources monhmc	X	X	X	X
Remove the utilization data collected for a managed system from the HMC rmlparutil	X	X		X
Enable users to edit a text file on the HMC in a restricted mode rnvi	X	X	X	X
Restore hardware resources after a DLPAR failure rsthwres		X		
Restore upgrade data on the HMC rstupgdata	X	X		X

Table 9. Command line tasks, associated commands, and user roles (continued)

Command line tasks	User roles/IDs			
	Operator (hmcoperator)	**Super Administrator (hmcsuperadmin)**	**Viewer (hmcviewer)**	**Service Representative (hmcservicerep)**
Transfer a file from the HMC to a remote system sendfile	X	X	X	X
chsvc	X	X		X
lssvc	X	X	X	X
chstat	X	X		X
lsstat	X	X	X	X
chpwdpolicy		X		
lspwdpolicy	X	X	X	X
mkpwdpolicy		X		
rmpwdpolicy		X		
expdata		X		

Session handling

Learn about session limitations in the Hardware Management Console (HMC).

Session limitations

The HMC does not support disconnected sessions. A session logoff and a session disconnect are both considered as a session logoff. This means that you cannot reconnect to the same session to resume your task or tasks that were initiated from a previous session. Every login through the HMC creates a new session.

1. If you initiate long running tasks from the HMC interface and then log off from the session, the long running tasks continue to run in the background. However, when you log in again, a new session is created and the task progress panels (which helps track the progress of the previous tasks) are no longer available. In this scenario, if you need to check the progress of the tasks that were initiated from a previous session, you can run the respective command line interface (CLI) commands, check the state of the managed resource, or check the console event logs.

   Note: Some examples of long running tasks include the following tasks:

   System management for servers:

   - Deploy system plan
   - Code update
   - Hardware - Prepare for hot repair or upgrade

   System management for partitions:

- DLPAR memory in large units in the order of Terabytes
- Live Partition Mobility (LPM)
- Suspend or resume

HMC management:

- Backup management console data
- Restore management console data
- Save upgrade data

2. If you fail to reauthenticate within the time that is specified in the verify timeout settings, you are automatically logged off from the current session.
3. The idle timeout user property task is not functional. The HMC interface uses the default value of **0** for the idle timeout setting. If you set a different value for this setting, it is ignored.

   Note: Session, idle, and verify timeout properties are set for a user and it can be different for different users on the same HMC.

Version mismatch state for a managed system

The **Version mismatch** state can occur when the redundant or dual Hardware Management Consoles (HMCs) that manage the same server are at different version and release levels.

The **Version mismatch** state can occur for any of the following reasons:

- FSP firmware and HMC versions are incompatible.
- An HMC Version 7.7.8 or later is connected to a server that was managed by a newer version of the HMC.
- An HMC Version 7.7.8 or later is connected to a server that was managed by a lower version of the HMC and does not have enough space present to upgrade the data to HMC Version 7.7.8 or later.
- The hypervisor or server brand or model is not supported by this version of the HMC.

To recover from the **Version mismatch** state, select the appropriate action, depending on the reference code that is displayed:

- **Save Area Version Mismatch**

  HMC Version 7.7.8 and later blocks attempts to manage a server with a configuration at a newer level by posting a new **Connection error** state and reference code. If an HMC Version 7.7.8 or later is connected to a server that was managed by a newer version of the HMC that updated the configuration format, then the HMC reports a connection error of **Version mismatch** with the reference code **Save Area Version Mismatch**. This error prevents accidental corruption of the configuration.

  If you want to continue on a lower HMC version, then you must first initialize the server in the lower version of the HMC before you proceed to run any operation.

- **Profile Data Save Area is full**

  The HMC uses a storage area on each managed server to store the server configuration, primarily PowerVM partition profiles. HMC Version 7.8.0 and later increases the usage of the storage area by adding another (mostly hidden) profile for each partition. Servers that already contain many profiles might not have sufficient space to allow the HMC Version 7.8.0 and later to run properly.

  HMC Version 7.8.0 and later checks for sufficient space in this storage area and stops the connection process with a connection state of **Version mismatch** and a reference code of **Profile Data Save Area is full** if sufficient space does not exist.

- **Connecting 0000-0000-00000000 (Unsupported Hypervisor)**

  A connection state of **Version mismatch** and a reference code of **Connecting 0000-0000-00000000 (Unsupported Hypervisor)** is returned when the server is configured for a hypervisor other than PowerVM.

To recover from this state, first start the ASM by selecting the server with the **Version mismatch** and selecting **Operations** and then **Launch Advanced System Manager (ASM)**.

On models that support multiple hypervisors, the hypervisor mode setting can be found in the ASM by selecting **System Configuration** and then **Hypervisor Configuration**. The hypervisor mode shows a setting of either PowerVM or OPAL.

If OPAL is the wanted configuration, then you must remove this connection from the HMC by selecting **Connections** and then **Reset or Remove Connections**. Next, select **Remove Connections** and click **OK**.

Note: The OPAL hypervisor is not supported on the HMC.

If PowerVM is the wanted configuration, select **PowerVM** from the hypervisor mode menu and click **Continue**.

Note: The setting can be changed only when the server is powered off. To power off the server select **Power/Restart Control** and then **Power On/Off System**. Click **Save Settings and Power off** .

- **Connection not allowed**

  A connection state of **Version mismatch** and a reference code of **Connection not allowed 0009-0008-00000000** is returned when the FSP firmware and HMC versions are incompatible.

  To recover from this state, install an HMC version that supports the managed server model.

For more information about correction a **Version mismatch** state, see Version mismatch errors.

Systems Management for Servers

Systems Management displays tasks to manage servers and logical partitions. Use these tasks to set up, configure, view current status, troubleshoot, and apply solutions for servers.

These tasks are listed when a managed system is selected. The tasks listed in the menu pod change as selections are made in the work area.

System content pane

View and monitor the state, health, and capacity information of all the systems that are connected to the management console.

The content pane, in the middle portion of the window, displays all the available systems and the associated information for each server. You can choose to display the information in a table view or a gallery view.

Each system displays the current state of the system, the number of central processing units (CPUs) that are in use, CPUs that are available, the amount of random access memory (RAM) that is in use, and the RAM that is available. Additionally, if you choose a tabular format view, you can filter the information to be displayed by selecting the drop-down arrow in the upper-right corner of the table. Select the check box corresponding to the field that you want to display on the table. If you are using HMC Version 9.1.930, in the **All Systems** table, you can also view information about the activated and deferred firmware levels.

You can click the **properties** icon to display the following information:

- Current state
- Reference code
- Machine type
- Serial number
- System location
- Firmware level
- Group tags
- Attention LED

You can click the **capacity** icon to display the following information:

- Date of collection.
- Processor usage (installed, activated, available, average, and peak). The bar graph and the numerical value show the average usage (average divided by activated) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by activated). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Memory allocation (installed, activated, available, average, and peak). The bar graph and the numerical value show the average usage (average divided by activated) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by activated). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Network I/O usage (sent and received in kilobytes per second and in packets per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Storage I/O usage (written and read information in kilobytes per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Data collection.

You can hover over the systems in the **All systems** window to view the system model description.

Operations

Operations contains the tasks for operating managed systems.

Power Off

Shut down the managed system. Powering off the managed system will make all partitions unavailable until the system is again powered on.

Before you power off the managed system, ensure that all logical partitions have been shut down and that their states have changed from `Running` to `Not Activated`. For more information on shutting down a logical partition, see "Shut Down" on page 58

If you do not shut down all logical partitions on the managed system before you power off the managed system, the managed system shuts down each logical partition before the managed system itself powers off. This can cause a substantial delay in powering off the managed system, particularly if the logical partitions are not responsive. Further, the logical partitions might shut down abnormally, which could result in data loss and further delays when you activate the logical partitions once more.

Choose from the following options:

Normal power off

The Normal power off mode shuts down the system's operations in a controlled manner. During the shutdown, programs running active jobs are allowed to perform cleanup (end-of-job processing).

Fast power off

The Fast power off mode shuts down the system by stopping all active jobs immediately. The programs running those jobs are not allowed to perform any cleanup. Use this option when you need to shut down the system because of an urgent or critical situation.

Power On

Use the **Power On** task to start a managed system.

Choose from the following options to power on your managed system:

Normal: Select this option to specify that the HMC uses the current setting for the partition start policy to determine how to power on the managed system. The current setting can be one of the following values:

- **Auto-Start Always**: This option specifies that the HMC power on logical partitions automatically after the managed system powers on. If powering on the managed system is the result of a user action, the HMC starts all partitions that are configured for automatic startup. If powering on the managed system is the result of an automatic recovery process, the HMC starts only those logical partitions that were running at the time the system is powered off. This option is always available for selection.
- **Stop at Partition Standby**: This option specifies that logical partition startup is in standby mode after the managed system powers on and the HMC does not start any logical partitions when the managed system powers on. If powering on the managed system is the result of an automatic recovery process and the HMC is used to start a logical partition, the HMC starts all logical partitions that were running at the time the system is powered off. This option is available for selection only when the firmware for the managed system does not support advanced IPL capabilities.
- **Auto-Start for Auto-Recovery**: This option specifies that the HMC power on logical partitions automatically only after the managed system powers on as the result of an automatic recovery process. This option is available for selection only when the firmware for the managed system supports this advanced IPL capability.
- **User-Initiated**: This option specifies that the HMC does not start any logical partitions when the managed system powers on. You must start logical partitions manually on the managed system by using the HMC. This option is available for selection only when the firmware for the managed system supports this advanced IPL capability.

You can set the partition start policy from the **Power On Parameters** page of the **Properties** task for the managed system.

System profile: Selecting this power-on option specifies that the HMC power on the system and its logical partitions based on a predefined system profile. When you select this power-on option, you must select the partition profile that you want the HMC to use to activate logical partitions on the managed system.

Hardware Discovery: Selecting this power-on option specifies that the HMC run the hardware discovery process when the managed system powers on. The hardware discovery process captures information about all I/O devices, in particular those devices that are not currently assigned to partitions. When you select the hardware discovery **power on** option for a managed system, the managed system is powered on into a special mode that performs the hardware discovery. After the Hardware Discovery process is complete, the system will be in Operating state with any partitions in the power-off state. The hardware discovery process records the hardware inventory in a cache on the managed system. The collected information is then available for use when you display data for I/O devices or when you create a system plan based on the managed system. This option is available only if the system is capable of using the hardware discovery process to capture I/O hardware inventory for the managed system.

Power Management

You can reduce the processor power consumption of the managed system by enabling power saver mode.

About this task

To enable power saver mode, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to enable the power saver mode and click **Actions** > **View All Actions**.
3. Select **Power Management** under **Operations**.
4. Choose from any of the following Power Saver mode options:
   - **Static**: Reduces the power consumption by reducing the processor clock frequency and the voltage to fixed values. This option delivers predictable performance while reducing the power consumption.

- **Dynamic Power Saver mode**: Delivers power savings by varying the processor frequency and voltage that is based on the utilization of the system processors. In Dynamic Power Saver Mode, the system firmware balances performance and power consumption.
- **Maximum Performance mode**: Causes the processor frequency to be set at a specified fixed value. You can set the maximum limit of the processor frequency and power consumption of the system.

Note: If you enable a Power Saver mode, processor frequencies, processor usage will change, and power consumption will change. It also causes varying performance.

5. You can also choose to enable or disable the **Idle Power Saver** mode. When it is enabled, it reduces the energy consumption when the system is in idle state.

   Note: Setting the Power Saver and Idle Power Saver modes are independent operations.

Schedule Operations

Create a schedule for certain operations to be performed on the managed system without operator assistance.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule power on or off operations for a managed system.

The Scheduled Operations task displays the following information for each operation:

- The processor that is the object of the operation.
- The scheduled date
- The scheduled time
- The operation
- The number of remaining repetitions

From the **Scheduled Operations** window, you can perform the following tasks:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

You can schedule an operation to occur once or you can schedule it to repeat. You must provide the time and date that you want the operation to occur. If you want the operation to repeat, you are asked to select the following options:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)
- The total number of repetitions. (required)

The operations that you can schedule for the managed system include the following operations:

Activate on a System Profile
Schedules an operation on a selected system for scheduling activation of a selected system profile.

Backup Profile Data
Schedules an operation to back up profile data for a managed system.

Power Off Managed System
Schedules an operation for a system power off at regular intervals for a managed system.

Power On Managed System
Schedules an operation for a system power-on at regular intervals for a managed system.

Manage Utility CoD processors
Schedules an operation for managing how your Utility CoD processors are used.

Manage Utility CoD processor minute usage limit
Creates a limit for Utility CoD processor usage.

Modify a Shared Processor Pool
Schedules an operation for modifying a shared processor pool.

Move a partition to a different pool
Schedules an operation for moving a partition to a different processor pool.

Change power saver mode on a managed system
Schedules an operation for changing a managed system's power saver mode.

Monitor/Perform Dynamic Platform Optimize
Schedules an operation for performing dynamic platform optimization and for sending an email notification alert to a user.

To schedule operations on the managed system, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select one or more managed systems and click **Actions** > **Schedule Operations**.
3. From the **Scheduled Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, click **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range...**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
4. To close the window, click **Options** and then click **Exit**.

Launch ASM Interface

The Hardware Management Console (HMC) can connect directly to the Advanced System Management Interface (ASMI) for a selected system.

The ASMI is an interface to the service processor that allows you to manage the operation of the server, such as auto power restart, and to view information about the server, such as the error log and vital product data.

To connect to the Advanced System Management Interface, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. In the content area, select one or more managed systems and click **Actions** > **View All Actions** > **Launch Advanced System Management (ASM)**.

Rebuild

You can extract the configuration information from the managed system and rebuild the information on the Hardware Management Console (HMC).

This task does not disrupt the operation of the running server.

Rebuilding the managed system updates the information on the HMC about the managed system. Rebuilding the managed system is useful when the state of the managed system is `Incomplete`. The Incomplete state means that the HMC cannot gather complete information from the managed system about logical partitions, profiles, or resources.

Rebuilding the managed system is different from refreshing the **HMC** window. When the managed system is rebuilt, the HMC extracts the information from the managed system. You cannot start other tasks while the HMC rebuilds the managed system. This process can take several minutes.

Change Password

Change the Hardware Management Console (HMC) access password on the selected managed system.

After the password is changed, you must update the HMC access password for all other HMCs from which you want to access this managed system.

Enter the current password and then, enter a new password and verify it by entering it again.

Attention LED

View system attention LED information, light specific LEDs to identify a system component, and test all LEDs on a managed system.

The system provides several LEDs that help identify various components, such as enclosures or field replaceable units (FRUs), in the system. For this reason, they are called **Identify** LEDs. Individual LEDs are on or near the components. The LEDs are located either on the component itself or on the carrier of the component (for example, memory card, fan, memory module, or processor). LEDs are either green or amber. Green LEDs indicate either of the following states:

- Electrical power is present.
- Activity is occurring on a link. (The system might be sending or receiving information).

Amber LEDs indicate a fault or identify condition. If your system or one of the components on your system has an amber LED turned on or flashing, identify the problem and take the appropriate action to restore the system back to normal.

You can activate or deactivate the following types of identify LEDs:

Identify LED for an enclosure
If you want to add an adapter to a specific drawer (enclosure), you need to know the machine type, model, and serial number (MTMS) of the drawer. To determine whether you have the correct MTMS for the drawer that needs the new adapter, you can activate the LED for a drawer and verify that the MTMS corresponds to the drawer that requires the new adapter.

Identify LED for a FRU associated with a specified enclosure
If you want to attach a cable to a specific I/O adapter, you can activate the LED for the adapter that is a field replaceable unit (FRU), and then physically verify where to attach the cable. This step can be especially useful when you have several adapters with open ports.

You can deactivate a system attention LED or a logical partition LED. For example, you might determine that a problem is not a high priority and decide to repair the problem later. However, you want to be alerted if another problem occurs, so you must deactivate the system attention LED so that it can be activated again if another problem occurs.

Choose from the following options:

Turn Attention LED Off
From this task, you can deactivate the system attention LED.

Identify Attention LED
Displays the current Identify LED states for all the location codes that are contained in the selected enclosure. From this task, you can select a single location code or multiple location codes to operate against and activate or deactivate one or more LEDs by selecting the corresponding button.

Test Attention LED

Initiates an LED Lamp Test against the selected system. All LEDs activate for several minutes.

Connections

You can view the Hardware Management Console (HMC) connection status to service processors or frames, reset those connections, connect another HMC to the selected managed system, or disconnect another HMC.

If you select a managed system in the work area, the following tasks pertain to that managed system. If you select a frame, the tasks pertain to that frame.

Service Processor Status

View information about the status of the Hardware Management Console (HMC) connection to the service processors on the managed system.

About this task

To show the service processor connection status to the service processors on the managed system, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view the service processor connection status and click **Actions** > **View All Actions** > **Service Processor Status**.

Reset or Remove Connections

Reset or remove a managed system from the Hardware Management Console (HMC) interface.

About this task

To reset or remove connections, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server that you want to reset or remove and click **Actions** > **Reset or Remove System Connection**.
3. Select one of the options from **Reset Connection** or **Remove Connection** and click **OK**.

Disconnect Another HMC

You can disconnect a connection between a selected Hardware Management Console (HMC) and the managed server.

About this task

To disconnect another HMC, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to disconnect another Management Console and click **Actions** > **View All Actions** > **Disconnect Another HMC**.
3. Select an HMC from the list and click **OK**.

System Templates

System templates contain configuration details for resources such as system properties, shared processor pools, reserved storage pool, shared memory pool, Host Ethernet adapters, and SR-IOV adapters. Many of the system settings that you previously configured by using separate tasks are available in the **Deploy System from Template** wizard. For example, you can configure the Virtual I/O Servers, virtual network bridges, and virtual storage settings when you use the wizard to deploy a system from a system template.

The template library includes predefined system templates, which contain configuration settings based on common usage scenarios. Predefined system templates are available for your immediate use.

You can also create custom system templates that contain configuration settings that are specific to your environment. You can create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration.

Deploy System from Template

You can deploy systems by using system templates that are available in the template library of the Hardware Management Console (HMC). The Deploy System from Template wizard guides you to provide target system-specific information that is required to complete the deployment of the selected system.

Create Partition from Template

You can create a partition by using partition templates that are available in the template library of the Hardware Management Console (HMC). The Create a Partition from Template wizard guides you through the deployment process and configuration steps.

Capture Configuration as Template

You can capture the configuration details of a running server and save the information as a custom system template by using the Hardware Management Console (HMC). This function is useful if you want to deploy multiple servers with the same configuration. If you want to use a predefined template, you do not need to complete this task.

To capture configuration as a template, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **View All Actions**
3. Click **Capture Configuration as Template with Physical I/O** or **Capture Configuration as Template without Physical I/O**.
4. Enter a template name and description, and then click **OK**.

Use the online Help if you need additional information about capturing the configuration as a template.

Updates

Display tasks to view system information, manage updates on your Hardware Management Console (HMC), or check system readiness.

View System Information

Display information on a selected system from the Hardware Management Console (HMC).

To view the network topology, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates** > **View System Information**.

3. Select a LIC repository from the list and click **OK**.
4. When you have completed this task, click **Close**.

Use the online Help if you need additional information for viewing system information of the HMC.

Change Licensed Internal Code

Change the Licensed Internal Code of a managed system by using your Hardware Management Console (HMC).

You can change the Licensed Internal Code for the current release or to a new release.

To change the Licensed Internal Code, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates**.
3. Select **Change Licensed Internal Code** > **for the Current Release** or **Change Licensed Internal Code** > **to a New Release**.

   Note: Click **Start Change Licensed Internal Code** wizard to start a guided update of managed system, power, and I/O Licensed Internal Code (LIC). Click **View System Information** to examine current LIC levels, including retrievable levels. Click **Select Advanced Features** to update the managed system and power the LIC with more options and more targeting choices.
4. Select an action from the list and click **OK**.
5. When you complete this task, click **Close**.

Use the online Help if you need additional information for changing the Licensed Internal Code of the HMC.

Check System Readiness

Check the readiness of the Licensed Internal Code of a selected system from the Hardware Management Console (HMC).

To check system readiness, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates** > **Check System Readiness**.
3. When you have completed this task, click **OK**.

Use the online Help if you need additional information for checking system readiness of the HMC.

SR-IOV Firmware Update

Update the driver firmware for SR-IOV adapters on your Hardware Management Console (HMC).

Note: The adapter must be in shared mode.

To update the firmware for SR-IOV adapters, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view the system information and click **Actions** > **Updates** > **SR-IOV Firmware Update**.
3. Select and right-click an adapter or adapters to get the context menu.
4. Select the type of firmware update to start.

   Note: Either the adapter driver firmware can be updated or both the adapter driver and adapter firmware can be updated. During the update operation of the adapter or adapter driver firmware, configured logical ports on the adapter might experience a temporary disruption of network traffic. Each adapter can take between 2 - 5 minutes to update. Updates are performed serially.

5. When you have completed this task, click **Close**.

Use the online Help if you need additional information for updating the driver or firmware for SR-IOV adapters.

Legacy

You can view **legacy** tasks that are available on the Hardware Management Console (HMC).

If you select a managed system in the work area, the following **legacy** tasks pertain to that managed system.

Partition Availability Priority

Use this task to specify the partition-availability priority of each logical partition on this managed system.

The managed system uses partition-availability priorities when a processor fails. If a processor fails on a logical partition and unassigned processors are not available on the managed system, then the logical partition can acquire a replacement processor from logical partitions with a lower partition-availability priority. This task allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

You can change the partition availability priority for a partition by selecting a partition and by choosing an availability priority from the list.

Use the online Help if you need additional information about prioritizing partitions.

View Workload Management Groups

Display a detailed view of the workload management groups that you specify for the managed system.

Each group displays the total number of processors, processing units for partitions that use shared mode processing, and the total amount of memory that is allocated to the partitions in the group.

Manage System Profiles

A system profile is an ordered list of partition profiles that is used by the Hardware Management Console (HMC) to start the logical partitions on a managed system in a specific configuration.

When you activate the system profile, the managed system attempts to activate each partition profile in the system profile in the order specified. A system profile helps you to activate or change the managed system from one set of logical partition configurations to another.

You can create a system profile that has a partition profile with overcommitted resources. You can use the HMC to validate the system profile against the currently available system resources and against the total system resources. Validating your system profile ensures that your I/O devices and processing resources are not overcommitted, and it increases the likelihood that the system profile can be activated. The validation process estimates the amount of memory that is needed to activate all of the partition profiles in the system profile. A system profile can pass validation and yet not have enough memory to be activated.

Use this task to complete the following tasks:

- Create new system profiles.
- Create a copy of a system profile.
- Validate the resources that are specified in the system profile against the resources available on the managed system. The validation process indicates whether any of the logical partitions in the system profile are already active and whether the uncommitted resources on the managed system can meet the minimum resources that are specified in the partition profile.
- View the properties of a system profile. From this task, you can view or change an existing system profile.
- Delete a system profile.

- Activate a system profile. When you activate a system profile, the managed system attempts to activate the partition profiles in the order that is specified in the system profile.

Use the online Help if you need additional information about managing system profiles.

Manage Partition Data

A partition profile is a record on the HMC that specifies a possible configuration for a logical partition. When you activate a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

A partition profile specifies the wanted system resources for the logical partition and the minimum and maximum amounts of system resources that the logical partition can have. The system resources that are specified within a partition profile includes processors, memory, and I/O resources. The partition profile can also specify certain operating settings for the logical partition. For example, you can set a partition profile such that, when the partition profile is activated, the logical partition is set to start automatically the next time that you power on the managed system.

Each logical partition on a managed system that is managed by an HMC has at least one partition profile. You can create more partition profiles with different resource specifications for your logical partition. If you create multiple partition profiles, you can designate any partition profile on the logical partition to be the default partition profile. The HMC activates the default profile if you do not select a specific partition profile to be activated. Only one partition profile can be active at one time. To activate another partition profile for a logical partition, you must shut down the logical partition before you activate the other partition profile.

A partition profile is identified by partition ID and profile name. Partition IDs are whole numbers that are used to identify each logical partition that you create on a managed system, and profile names identify the partition profiles that you create for each logical partition. Each partition profile on a logical partition must have a unique profile name, but you can use a profile name for different logical partitions on a single managed system. For example, logical partition 1 cannot have more than one partition profile with a profile name of normal, but you can create a profile named normal for each logical partition on the managed system.

When you create a partition profile, the HMC shows you all of the resources available on your system. The HMC does not verify whether another partition profile is using a portion of these resources. Therefore, it is possible for you to overcommit resources. When you activate a profile, the system attempts to allocate the resources that you assigned to the profile. If you overcommit resources, the partition profile is not activated.

For example, you have four processors on your managed system. Partition 1 profile A has three processors, and partition 2 profile B has two processors. If you attempt to activate both of these partition profiles at the same time, partition 2 profile B fails to activate because you overcommitted processor resources.

When you shut down a logical partition and reactivate the logical partition by using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition by using dynamic logical partitioning are lost when you reactivate the logical partition that uses a partition profile. This action is required when you want to undo dynamic logical partitioning changes for the logical partition. However, this action is not required if you want to reactivate the logical partition that uses the resource specifications that the logical partition had when you shut down the managed system. Therefore, keep your partition profiles up to date with the latest resource specifications. You can save the current configuration of the logical partition as a partition profile. This task avoids having to change partition profiles manually.

If you shut down a logical partition whose partition profiles are not up to date, and if the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system by using the partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

Use the Manage Partition Data tasks to complete the following tasks:

- Restore partition data. If you lose partition profile data, use the restore task in one of the following ways:
  - Restore partition data from a backup file. Profile modifications that are completed after the selected backup file was created are lost.
  - Restore merged data from your backup file and recent profile activity. The data in the backup file takes priority over recent profile activity if the information conflicts.
  - Restore merged data from recent profile activity and your backup file. The data from recent profile activity takes priority over your backup file if the information conflicts.
- Initialize partition data. Initializing the partition data for a managed system deletes all of the currently defined system profiles, partitions, and partition profiles.
- Back up a partition profile to a file.
- Back up partition data to a file.

Use the online Help if you need additional information about managing partition data.

Utilization Data

You can set the Hardware Management Console (HMC) to collect resource utilization data for a specific managed system or for all systems the HMC manages.

The HMC collects utilization data for memory and processor resources. You can use this data to analyze trends and make resource adjustments. The data is collected into records that are called events. Events are created at the following times:

- At periodic intervals (30 seconds, 1 minute, 5 minutes, 30 minutes, hourly, daily, and monthly).
- When you make system-level and partition-level state and configuration changes that affect resource utilization.
- When you start, shut down, and change the local time on the HMC.

You must set the HMC to collect utilization data for a managed system before utilization data can display for the managed system.

Use the **Change Sampling Rate** task to enable, set and change the sampling rate, or to disable sampling collection.

Create Partition

You can quickly create partitions with minimum resources.

To create a partition, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to create a partition and click **Actions** > **View System Partitions**.
3. Click **Create Partition**.
4. Complete the required information in the **Basic Partition Configuration**, **Processor Configuration**, and **Memory Configuration** tabs. If you want to assign all the system resources to the partition, select the **Assign all system resources** check box.
5. To create multiple partitions, move the slider to the right and select the **Multiple Partitions View**.
6. To add a new partition definition, click the **(+)** sign located on the top of the partition table.
7. Select the added partition and complete the required information in the **Basic Partition Configuration**, **Processor Configuration**, and **Memory Configuration** tabs. In the **Basic Partition Configuration** tab, you can provide details about the number of partition instances that you want to create. You can create a maximum of 20 partition instances.
8. To remove an existing partition, select the partition that you want to remove and click the **(-)** sign.

9. Click **OK**.

Use the online Help if you need additional information about this task.

Note: If the managed system supports virtual serial number and if the managed system is not in an Enterprise Pool 2.0, the **Virtual Serial Number** can be specified in the **Basic Partition Configuration** tab.

When the firmware level is at FW950 and the managed system has logical partitions that are assigned with virtual serial numbers, the managed system cannot be added to an Enterprise Pool 2.0. Also, if the managed system is in an Enterprise Pool 2.0, virtual serial number cannot be assigned to the logical partitions.

Properties

Displays the properties of the selected managed system. This information is useful in system and partition planning and resource allocation.

To open the properties tasks that are available for your system, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Properties** and then select the properties task that you want to perform from the list.

General Settings

View or change the general and advanced settings for the managed system.

These properties include the following tabs:

General Properties

The **General Properties** tab displays the system's name, serial number, model and type, state, attention led state, service processor version, maximum number of partitions, assigned service partition (if designated), and power off policy information.

Migration

View the partition mobility properties and change the migration policy for inactive partitions on the managed system.

Power-On Parameters

From the **Power-On Parameters** tab, you can change the power-on parameters for the next restart by changing the values in the **Next Value** fields. These changes are only valid for the next managed system restart.

Advanced

The **Advanced** tab displays huge page memory capabilities on the managed system, including available huge page memory, configurable huge page memory, current page size, and current maximum huge page memory. To change memory allocation on systems with huge page table support, set the Requested huge page memory (in pages) field to the wanted memory. To change the requested value for huge page memory, the system must be powered off.

The **Barrier Synchronization Register (BSR)** option displays array information.

The **Processor Performance** option displays the TurboCore mode and the System Partition Processor Limit (SPPL). You can set the next TurboCore mode and the next SPPL value. The SPPL applies to both dedicated processor partitions and shared processor partitions.

The **Memory Mirroring** option displays the current mirroring mode and the current system firmware mirroring status. You can set the next mirroring mode. You can also start the memory optimization tool.

You can view the VTPM settings.

Processor, Memory, I/O

View or change the memory, processor, and physical I/O resource settings for the managed system.

These properties include the following tabs:

Processor

The **Processor** tab displays information about the processors of the managed system, which includes:

- installed processing units
- unconfigured processing units
- available processor units
- available with stealable processor units
- configurable processing units
- minimum number of processing units per virtual processor
- maximum number of shared processor pools

The **Available with stealable** field displays the information about the available processing units, which is the sum of the available processing units in the managed system and the number of stealable processing units.

The stealable processor units value is the sum of the processor resources that are assigned to all the powered off or hibernated partitions on the managed system.

Notes:

- The information about stealable processor units is available only when the managed system is in the standby state or in the operating state.
- If the managed system is licensed with Power® IFL processor and if the firmware level is at FW910, or later, the **Available (with stealable)** field is displayed.
- When a Power10 system is licensed with some IFL processors, the tab also displays the information about the remaining processors that are available for running the AIX or IBM i partitions.

Memory

The **Memory** tab displays information about the memory of the managed system, which includes:

- installed memory
- unconfigured memory
- available memory
- available with stealable memory
- configurable memory
- memory region size with `Current Value` and `Next Value` fields that specify the current and next available value.

  Note: You can change the size of the Logical Memory Block (LMB) by changing the values in the `Next Value` field. The changes of this field are applied only after the next managed system restart.
- current memory available for partition usage
- system firmware current memory

The **Available with stealable** field displays the information about the available memory, which is the sum of available memory in the managed system and the amount of stealable memory resources. The tab also displays the maximum number of memory pools that are available.

Note: The information about stealable memory resources is available only when the managed system is in the standby state or in the operating state.

Physical I/O adapters

The **Physical I/O Adapters** tab displays the physical I/O resources for the managed system. The assignment of I/O slots and partition, the adapter-type, and the slot LP limit information are displayed. The physical I/O resources information is grouped by units.

- The **Adapter Description** column displays the physical description of each resource.
- The **Physical Location Code** column displays the physical location code of each resource.
- The **Owner** column displays who currently owns the physical I/O. The value of this column can be any of the following values:
  - When a single root I/O virtualization (SR-IOV) adapter is in the shared mode, **Hypervisor** is displayed in this column.
  - When an SR-IOV adapter is in the dedicated mode, **Unassigned** is displayed when the adapter is not assigned to any partition as a dedicated physical I/O.
  - When an SR-IOV adapter is in the dedicated mode, the logical partition name is displayed when the adapter is assigned to any logical partition as a dedicated physical I/O.
- The **Bus Number** column displays the bus number of the resource.
- The **I/O Pools** button displays all of the I/O pools found in the system and the partitions that are participating in the pools.

Persistent Memory

The **Persistent Memory** tab displays details about the Virtual Persistent Memory and Hybrid Memory Subsystem (devices or volumes) of the managed system.

Hybrid Memory Subsystem (HMS)

The **HMS** pane displays all the persistent memory devices and persistent memory volumes of the managed system.

The **Persistent Memory Devices** table displays information about the persistent memory devices, which includes:

- **Location code**: Displays the location code of the persistent memory device.
- **Status**: Indicates whether the persistent memory device is functional.
- **Device**: Displays the device unique ID of the persistent memory device.
- **Block size**: Displays the block size of the persistent memory device.
- **Available size**: Displays the available size of the persistent memory device.
- **Total size**: Displays the total size of the persistent memory device.
- **Erase Capable**: Indicates whether the HMS device supports initiating the Erase operation for the persistent memory device data.
- **Serial number**: Displays the serial number of the persistent memory device.
- **WWID**: Displays the worldwide ID of the persistent memory device.
- **Current volumes**: Displays the current size of the HMS volume that is allocated from the HMS device.
- **Maximum volumes**: Displays the maximum volume size that can be allocated to the persistent memory device.

You can perform the following operations on a persistent memory device of the managed system:

- To create volume for a persistent memory device, select a persistent memory device from the **Persistent Memory Devices** table and then click **Action** > **Create Volume**.
- To format a persistent memory device, select a persistent memory device from the **Persistent Memory Devices** table and then click **Action** > **Format**.
- To format and restore a persistent memory device, select a persistent memory device from the **Persistent Memory Devices** table and then click **Action** > **Format and Restore**.

- To initiate erase for a persistent memory device, select a persistent memory device displayed in the **Persistent Memory Devices** table and click **Action** > **Initiate Erase**.

The **Persistent Memory Volumes** table displays information about Persistent Memory (PMEM) volumes that are associated with the persistent memory device, which includes:

- **Name**: Displays the name of the persistent memory volume.
- **Size**: Displays the total size of the persistent memory volume.
- **Assigned partition**: Displays the name of the logical partition that is assigned to the persistent memory volume.
- **WWID**: Displays the worldwide ID of the persistent memory volume.
- **Current size**: Displays the current size of a HMS volume that is allocated from the HMS device.
- **Device**: Displays the device unique ID of the persistent memory device that is associated with the persistent memory volume
- **In use**: Indicates whether the persistent memory volume is active.

You can perform the following operations on a persistent memory volume that is associated with the persistent memory device:

- To modify a persistent memory volume, select a persistent memory volume from the **Persistent Memory Volumes** table and then click **Action** > **Modify**.
- To format a persistent memory volume, select a persistent memory volume from the **Persistent Memory Volumes** table and then click **Action** > **Format**.
- To delete a persistent memory volume, select a persistent memory volume from the **Persistent Memory Volumes** table and then click **Action** > **Delete**.

Virtual Persistent Memory

The **Virtual Persistent Memory** pane displays information about the persistent memory volumes, which includes:

- **Name**: Displays the name of the virtual persistent memory.
- **Size**: Displays the size of the virtual persistent memory.
- **Assigned partition**: Displays the name of the logical partition that is associated with the virtual persistent memory.
- **Current size**: Displays the size of the volume that is allocated to the virtual persistent memory from the total volume size.
- **Affinity**: Indicates whether the virtual persistent memory is internally divided into multiple smaller virtual persistent memory.

PowerVM

You can use the PowerVM function on the Hardware Management Console (HMC) to manage the system-level virtualization capabilities of your IBM Power Systems servers.

You can use the PowerVM task to manage virtual resources that are associated with a system, such as configuring a Virtual I/O Server (VIOS), virtual networks, and virtual storage. You can manage the PowerVM functions at the managed system level in response to changes in workloads or to enhance performance.

The PowerVM functions include the following tasks:

- Managing Virtual I/O Servers
- Managing virtual networks
- Managing virtual storage
- Managing hardware virtualized I/O (SR-IOV adapters)
- Managing a reserved processor pool
- Managing shared processor pools

- Managing a shared memory pool

Use the online Help if you need additional information about managing PowerVM.

Capacity on Demand

Activate disabled processors or memory that is installed on your managed server.

You can use Capacity on Demand (CoD) to nondisruptively activate (no startup needed) processors and memory. Capacity on Demand also gives you the option to temporarily activate capacity to meet intermittent performance needs, to activate extra capacity on a trial basis, and to access capacity to support operations in times of need.

Capacity on Demand Functions

Learn about the different Capacity on Demand functions that are available for your system.

You can use Capacity on Demand (CoD) to nondisruptively activate (no startup needed) processors and memory. Capacity on Demand also gives you the option to temporarily activate capacity to meet intermittent performance needs, to activate extra capacity on a trial basis, and to access capacity to support operations in times of need.

The **Capacity on Demand Processor** functions include the following tasks:

- View processor settings
- CUoD (permanent) processor
  - View CUoD code information
- On/Off processor
  - Manage
  - View billing information
  - View capacity settings
  - View code information
- Utility processor
  - Manage
  - View capacity settings
  - View code information
  - View shared processor utilization
- Trial processor
  - Stop trial
  - View capacity settings
  - View code information

The **Capacity on Demand Memory** functions include the following tasks:

- View memory settings
- CUoD (permanent) memory
  - View CUoD code information
- On/Off memory
  - Manage
  - View billing information
  - View capacity settings
  - View code information

- Trial memory
  - Stop trial
  - View capacity settings
  - View code information

Use the online Help if you need additional information about Capacity on Demand functions.

Licensed Capabilities

View and edit the runtime capabilities that are supported by the managed system.

You can view which licensed capabilities are active on your managed system. To activate a new licensed capability, click **Enter Activation Code** and enter the activation code.

The licensed functions that are available on the managed system include the following capabilities:

- Active Memory Sharing Capable
- Live Partition Mobility Capable
- Micro-Partitioning® Capable
- PowerVM Partition Simplified Remote Restart Capable
- SR-IOV Capable (Logical Port Limit)
- Virtual I/O Server Capable
- Active Memory Expansion Capable
- Active Memory Mirroring for Hypervisor Capable
- Coherent Accelerator Processor Interface (CAPI)
- AIX Enablement for 256-Core Partition Capable
- Dynamic Platform Optimization Capable
- IBM i 5250 Application Capable

Use the online Help if you need additional information about licensed capabilities.

Serviceability

Problem Analysis on the HMC automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. Use the **Serviceable Events Manager** task to view specific events for selected systems. However, if you notice a problem occurred or you suspect a problem is affecting the system but Problem Analysis has not reported it to you, use the **Create Serviceable Event** task to report the problem to your service provider.

To open the serviceability tasks that are available for your system, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the serviceability task that you want to perform from the list.

Tasks Log

View all the tasks that are currently running or completed on the Hardware Management Console (HMC).

To view the tasks log, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Tasks Log**.
2. You can view the following tabs in the tasks log:

- **Task name**: Displays the name of task.
- **Status**: Displays the current state of the task (running or completed).
- **Resource**: Displays the name of the resource.
- **Resource type**: Displays the type of resource.
- **Initiator**: Displays the name of the user that initiated the task.
- **Start time**: Displays the time that the task was initiated.
- **Duration**: Displays the amount of time that the task took to complete.

Use the online Help for additional information about viewing the tasks log.

Serviceability

Problem Analysis on the Hardware Management Console (HMC) automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. Use the **Serviceable Events Manager** task to view specific events for selected systems. However, if you notice a problem occurred or you suspect a problem is affecting the system but Problem Analysis has not reported it to you, use the **Create Serviceable Event** task to report the problem to your service provider.

To open the serviceability tasks that are available for your system, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the serviceability task that you want to perform from the list.

Serviceable Events Manager

Problems on your managed system are reported to the Hardware Management Console (HMC) as serviceable events. You can view the problem, manage problem data, call home the event to your service provider, or repair the problem.

To set the criteria for the serviceable events you to view, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceable events.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Click **Serviceable Events Manager**.
5. Provide event criteria, error criteria, and FRU criteria. If you do not want the results to be filtered, select **ALL**. Click **Refresh** to refresh the list of serviceable events based on the criteria filter values.

The **Serviceable Events Overview** window displays all of the events that match your criteria. The information that is displayed in the compact table view includes the following information:

- Problem Number
- PMH Number
- Reference Code - Click the Reference code to display a description of the problem reported and actions that can be taken to fix the problem.
- Status of the problem
- Last reported time of the problem
- Failing MTMS of the problem
- Originating HMC

The full table view includes more detailed information, including reporting partition ID, primary data event timestamp, duplicate count, notification type, first reported time, reporting name, reporting MTMS, firmware fix, and serviceable event text.

Select a serviceable event and use the **Action** menu to:

- **View Details**: Field-replaceable units (FRUs) associated with this event and their descriptions.
- **View Files**: View the files associated with the selected serviceable event.
- **View Reference Code Description**: View the description of the reference code associated with the selected serviceable event. The option will not be available if additional description is not available.
- **Call Home**: Report the event to your service provider.
- **Repair**: Start a guided repair procedure, if available.
- **Close Event**: After the problem is solved, add comments and close the event.
- **Add PMH Comment**: Add a comment to a PMH for a selected serviceable event. If a PMH number does not exist for a given problem, the Add PMH Comment option is not available.

Use the online Help if you need additional information on managing serviceable events.

Create Serviceable Event

This task reports problems that occurred on your Hardware Management Console (HMC) to the service provider (for example, the mouse does not work) or lets you test problem reporting.

Submitting a problem is dependent upon whether you have customized this Hardware Management Console to use the Remote Support Facility (RSF) and if it is authorized to automatically call for service. If so, the problem information and service request is sent to the service provider automatically with a modem transmission.

To report a problem on your Hardware Management Console, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Create Serviceable Event**.
3. From the **Create Serviceable Event** window, select a problem type from the list displayed.
4. Enter a brief description of your problem in the **Problem Description** input field and then click **Request Service**.

To test problem reporting from the **Report a Problem** window:

1. Select **Test automatic problem reporting** and enter *This is just a test* in the **Problem Description** input field.
2. Click **Request Service**. The problems are reported to the service provider for the Hardware Management Console. Reporting a problem sends to the service provider the information you provide on **Report a Problem** window, and machine information that identifies the console.

Use the online Help if you need additional information for reporting a problem or testing if problem reporting works.

Manage Dumps

Manage system, service processor, and power subsystem dumps for systems that are managed by the Hardware Management Console (HMC).

system dump
A collection of data from server hardware and firmware, either after a system failure or a manual request. Perform a system dump only under the direction of your next level of support or your service provider.

service processor dump
A collection of data from a service processor either after a failure, external reset, or manual request.

power subsystem dump

A collection of data from Bulk Power Control service processor. This process is only applicable to certain models of managed systems.

Use the **Manage Dump** task to complete the following tasks:

- Initiate a system dump, a service processor dump, or a power subsystem dump.
- Modify the dump capability parameters for a dump type before you initiate a dump.
- Delete a dump.
- Copy a dump to media.
- Copy a dump to another system by using file transfer protocol (FTP).
- Call home a dump by using the Call Home feature to transmit the dump back to your service provider, for example IBM Remote Support, for further analysis.
- View the offload status of a dump as it progresses.

Use the online Help if you need additional information for managing dumps.

Collect VPD

Copy Vital Product Data (VPD) to removable media.

The managed system has VPD that is stored internally. The VPD consists of information such as how much memory is installed, and how many processors are installed. These records can provide valuable information that can be used by remote service and service representatives so that they can help you keep the firmware and software on your managed system up to date.

Note: To collect VPD, you must have at least one operational partition. For more information, see Logical Partitioning.

The information in the VPD file can be used to complete the following types of orders for your managed system:

- Install or remove a sales feature.
- Upgrade or rollback a model.
- Upgrade or rollback a feature.

Using this task, this information can be sent to removable media (diskette or memory key) for use by you or your service provider.

Use the online Help if you need additional information for collecting VPD.

Type, Model, Feature

Edit or display the model, type, machine type model serial (MTMS), or configuration ID of an enclosure.

The MTMS value or configuration ID for an expansion unit might need to be edited during a replacement procedure.

Use the online Help if you need additional information for editing MTMS.

Hardware Operations

Add, exchange, or remove hardware from the managed system. Display a list of installed FRUs or enclosures and their locations. Select a FRU or an enclosure and launch a step-by-step procedure to add, exchange, or remove the unit.

To open the hardware tasks that are available for your system, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage hardware tasks.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Select the hardware operations task that you want to perform from the list.

Prepare for Hot Repair or Upgrade
Provides a summary of required actions to be performed to isolate a particular hardware component as part of a service procedure.

From the **Component List** table, you can select the component to be repaired using the location code on the system to be repaired as directed by an Authorized Service Representative.

Power On/Off Unit
Use the **Power On/Off Unit** task to power on or off an I/O unit.

Only units or slots that reside in a power domain can be powered on or off. The corresponding power on/off buttons are disabled for location codes that are not controllable by the HMC.

Add FRU
Locate and add a Field Replaceable Unit (FRU).

To add a FRU to a POWER10 system, complete the following steps:

1. Select an enclosure type from the **Enclosure** menu.
2. Select a FRU type from the displayed list of FRU types for this enclosure, and click **Next**.
3. Select a FRU location, then click **Next** to start the Add FRU procedure for the selected location.
4. Follow the **Procedure** details listed in the window. The **Operation** list displays the correct operations for the procedure that you want to execute. The procedures direct when to perform each operation. Wait until advised to perform the operation before selecting the operation in the **Operation** list. Click **Next** to continue to the next service procedure.

   Note: Select the FRU location image on the header to display the location of the part to be serviced. To see a larger version of the FRU location image in a separate window, click **Display FRU Location**.
5. Click **Finish** to end the service when you have completed the last service procedure.

When the managed system is a POWER8® or earlier, to add a FRU, complete the following steps:

1. Select an enclosure type from the **Add FRU** menu.
2. Select a FRU type from the menu.
3. Click **Next**.
4. Select a location code from the displayed menu.
5. Click **Add**.
6. Click **Launch Procedure**.
7. When you complete the FRU installation process, click **Finish**.

Exchange FRU
Use the **Exchange FRU** task to exchange one field replaceable unit (FRU) with another FRU.

When the managed system is a POWER10 or later, to exchange a FRU, complete the following steps:

1. Select an installed enclosure type from the **Enclosure** menu.
2. Select a FRU type to be replaced, from the displayed list of FRU types for this enclosure and click **Next**.
3. Select a installed FRU location, then click **Next** to start the Exchange / Replace FRU procedure for the selected FRU.
4. Follow the **Procedure** details listed in the window. The **Operation** list displays the correct operations for the procedure that you want to execute. The procedures direct when to perform each operation. Wait until advised to perform the operation before selecting the operation in the **Operation** list. Click **Next** to continue to the next service procedure.

   Note: Select the FRU location image on the header to display the location of the part to be serviced. To see a larger version of the FRU location image in a separate window, click **Display FRU Location**.
5. Click **Finish** when you complete the exchange procedure.

When the managed system is a POWER8 or earlier, to exchange a FRU, complete the following steps:

1. Select an installed enclosure type from the **Exchange FRU** menu.
2. From the displayed list of FRU types for this enclosure, select a FRU type.
3. Click **Next** to display a list of locations for the FRU type.
4. Select a location code for a specific FRU.
5. Click **Add** to add the FRU location to **Pending Actions**.
6. Select **Launch Procedure** to begin replacing the FRUs that are listed in **Pending Actions**.
7. Click **Finish** when you complete the installation.

Remove FRU
Use the **Remove FRU** task to remove a FRU from your managed system.

When the managed system is a POWER10 or later, to remove a FRU, complete the following steps:

1. Select an enclosure from the menu to display a list of FRU types that are currently installed in the selected enclosure.
2. Select a FRU type from the displayed list of FRU types available for removal from the selected system and click **Next**.
3. Select a FRU location, then click **Next** to start the Remove FRU procedure for the selected FRU .
4. Follow the **Procedure** details listed in the window. The **Operation** list displays the correct operations for the procedure you want to execute. The procedures direct when to perform each operation. Wait until advised to perform the operation before selecting the operation in the **Operation** list. Click **Next** to continue to the next service procedure.

   Note: Select the FRU location image on the header to display the location of the part to be serviced. To see a larger version of the FRU location image in a separate window, click **Display FRU Location**.
5. Click **Finish** when you complete the removal procedure.

When the managed system is a POWER8 or earlier, to remove a FRU, complete the following steps:

1. Select an enclosure from the menu to display a list FRU types that are currently installed in the selected enclosure.
2. From the displayed list of FRU types for this enclosure, select a FRU type.
3. Click **Next** to display a list of locations for the FRU type.
4. Select a location code for a specific FRU.
5. Click **Add** to add the FRU location to **Pending Actions**.
6. Select **Launch Procedure** to begin removing the FRUs listed in **Pending Actions**.
7. Click **Finish** when you complete the removal procedure.

Add Enclosure
Learn how to locate and add an enclosure.

To add an enclosure, complete the following steps:

1. Select an enclosure type, then click **Add**.
2. Click **Launch Procedure**.
3. When you complete the enclosure installation process, click **Finish**.

Remove Enclosure
Use the **Remove Enclosure** task to remove an enclosure.

To remove an enclosure, complete the following steps:

1. Select an enclosure type, then click **Add** to add the selected enclosure type's location code to **Pending Actions**.
2. Click **Launch Procedure** to begin removing the enclosures that are identified in **Pending Actions** from the selected system.

3. Click **Finish** when you complete the enclosure removal process.

Open MES
View MES order numbers and their states, for any MES operations active or inactive for the Hardware Management Console (HMC).

Use **Add MES Order Number** to add a new order number to the list. To add an order number, complete the following steps:

1. Click **Add MES Order Number**.
2. Enter new MES order number.
3. Click **OK**.

Close MES
Close MES order numbers.

Use **Close MES** to close a MES. To close a MES, complete the following steps:

1. Select an open MES order number from the table.
2. Click **OK**.

Setup FSP Failover
Set up a secondary service processor if your managed system's primary service processor fails.

FSP Failover is designed to reduce customer outages due to service processor hardware failures. If a redundant service processor is supported for the current system configuration, select **Setup** to set up FSP Failover for the selected managed system.

To set up the FSP failover, complete the following steps:

1. In the content pane under **FSP failover**, click **Setup**.
2. Click **OK** to enable automatic failover for the selected system.

Initiate FSP Failover
Initiate a secondary service processor if your managed system's primary service processor fails.

FSP Failover is designed to reduce customer outages due to service processor hardware failures. Select **Initiate** to start the FSP Failover for the selected managed system.

To start the FSP failover, complete the following steps:

1. In the content pane under **FSP failover**, click **Initiate**.
2. Click **OK** to start the automatic failover for the selected system.

Reference Code Log

Reference codes provide general diagnostic, troubleshooting, and debugging information.

View reference codes that are generated for the selected managed system. Reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem.

To view the reference code history, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **Reference Code Log**.
4. Select a specific reference code to view the details.

Use the online Help if you need additional information about this task.

RIO Configuration

View the current hardware topology and the last valid hardware topology.

Displays the current hardware and last valid hardware topology. Any discrepancies between the current topology and the last valid topology are identified as errors.

To view the hardware topology, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **RIO Configuration**.
4. View the hardware topology information.

Use the online Help if you need additional information about this task.

PCI Configuration

View information about the Peripheral Component Interconnect Express (PCIe) hardware topology.

The PCIe hardware topology utility provides information about the PCIe links that exist for each system.

To view the PCIe hardware topology, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Serviceability** and then click **PCI Configuration**.
4. View the PCIe hardware topology.

Use the online Help if you need additional information about this task.

Topology diagrams

Learn how to view the topology diagrams of a partition.

You can use the Hardware Management Console (HMC) to view the topology diagrams of a partition.

Viewing virtual networking diagrams

You can view the end-to-end network configuration for the selected system, by using the Hardware Management Console (HMC). The view of the virtual networks begins with the physical adapter cards and the physical ports that are connected to them. As you scroll down, you can see the defined virtual bridges, link aggregation devices, virtual switches, virtual networks, and partitions in the VIOS.

You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the network diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the end-to-end network configuration for the selected system by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Topology** and then click **Virtual Networking Diagram**.
4. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.

5. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
6. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols that are used in the virtual networking diagram.

Use the online Help if you need additional information about this task.

Viewing virtual storage diagrams

Two types of virtual storage diagrams are available; systems storage and partition storage. You can view the virtual storage configuration for the selected system, including the physical and virtual components of system storage, by using the HMC. You can also view the virtual storage configuration for a single partition in a particular system, including the physical and virtual components of storage assigned to that particular partition, by using the Hardware Management Console (HMC).

This diagram displays a high-level overview of the contents of the system or a single partition, and not specific component relationships. You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the storage diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the virtual storage configuration for the selected system or a single partition by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Topology** and then click **Virtual Storage Diagram**.

   Note: To view the virtual storage diagrams of a single partition in a particular system, select the partition of your choice and then expand **Topology** and click **Partition Virtual Storage Diagram**
4. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
5. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
6. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols that are used in the virtual storage diagram.

Use the online Help if you need additional information about this task.

Viewing SR-IOV and vNIC diagrams

You can view the SR-IOV and virtual Network Interface Controllers (vNIC) configuration for the selected system, including the physical and virtual components, by using the Hardware Management Console (HMC).

This diagram displays the relationships between SR-IOV adapters and other virtual components, such as vNIC. You can click a resource and drag to pan across the diagram. You can also double-click a resource to highlight that resource and the relationship between its various virtual and physical components in the network. To remove the highlighting, double-click in an empty area of the SR-IOV and vNIC diagram. To view more detailed information about a resource, you can right-click a resource and additional information is displayed in a click card. Alternatively, you can hover over the label of a resource area to display the name of the resource as a tooltip.

To view the end-to-end network configuration for the selected system by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceability tasks and click **Actions** > **View System Partitions**.
3. In the menu pod, expand **Topology** and then click **SR-IOV vNIC Diagram**.
4. Right-click a resource for the selected system to view more detailed information in a click card. You can also hover over the label of a resource area to display the name of the resource as a tooltip.
5. In the upper-right corner of the work pane, click the **zoom in** and **zoom out** icons to achieve the required level of magnification.

   Note: You can also zoom in and zoom out by using the scroll wheel on the mouse from within the diagram.
6. In the upper-right corner of the work pane, click the **Legend** icon to view an explanation of the symbols that are used in the SR-IOV and vNIC diagram.

Use the online Help if you need additional information about this task.

Systems Management for Partitions

Systems Management displays tasks that you can perform to manage servers and logical partitions. Use these tasks to set up, configure, view current status, troubleshoot, and apply solutions for partitions.

The following sets of tasks are represented when a partition is selected and is shown in the menu pod or content pane. The tasks that are listed in the menu pod change as selections are made in the work area.

Partition content pane

View and monitor the state, health, and capacity information of all the partitions that are connected to the management console.

The content pane, in the middle portion of the window, displays all the available partition and the associated information for each partition.

Each partition displays the current state of the partition, the reference code, the number of virtual processors that are allocated, and the amount of random access memory (RAM) that is allocated. Additionally, if you choose a tabular format view, you can filter the information to be displayed by selecting the drop-down arrow in the upper-right corner of the table. Select the check box corresponding to the field that you want to display on the table. If you are using HMC Version 9.1.930, or later, in the **All partition** table, you can also view the memory mode of all partitions that are associated with the managed system.

You can click the **properties** icon to display the following information:

- Current state
- System name
- Reference code
- Partition ID
- IP address
- Environment
- OS version
- RMC connection
- Last activated profile
- Contains physical I/O
- Group tags

- Attention LED

You can click the **capacity** icon to display the following information:

- Date of collection.
- Processor usage (type (dedicated, uncapped, or capped), entitled capacity, and virtual processors). When the processor type is dedicated, the bar graph and the numerical value show the used processor usage (used divided by assigned) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by assigned). When the processor type is uncapped, the bar graph and the numerical value show the used processor usage (used divided by the number of virtual processors) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by the number of virtual processors). When the processor type is capped, then the bar graph and the numerical value show the used processor usage (used divided by entitled) in a percentage. The threshold line on the bar graph shows the peak threshold (peak divided by entitled). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Memory allocation (allocated).
- Network I/O usage (sent and received in terabytes per second and in packets per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Storage I/O usage (written and read information in kilobytes per second). The bar graph and the numerical value show the average usage ((average transferred bytes divided by the number of adapters) divided by the maximum transferred bytes) in a percentage. The threshold line on the bar graph shows the high threshold (maximum transferred bytes). When the value exceeds the threshold, the bar graph changes to a dashed line format.
- Data collection.

Partition Properties

The **View Partition Properties** task displays the selected partition's properties. This information is useful in resource allocation and partition management. These properties include:

General

The **General** tab displays the partition's name, ID, environment, state, resource configuration, version of the operating system that is installed on the partition, boot mode to start the operating system, and the system on which the partition is located.

Notes:

- On Power10 processor-based systems, the image date of the OS that is installed on the partition and the expiration date of the AIX Update Access Key (UAK) for the managed system are displayed for AIX and AIX/Linux partitions.
- If the managed system supports virtual serial number and if the managed system is not in an Enterprise Pool 2.0, the Virtual Serial Number property is displayed in the **General** tab.

Click **Advanced settings** to also view the list of supported hardware accelerators for a logical partition and Quality of Service (QoS) credits for a specific hardware accelerator. This section is not displayed if the managed system does not support hardware accelerators.

Note: When the HMC is at V9.2.950.0, or later, and when the firmware is at level FW950, or later, the **KeyStore Size** value can be chosen in the range 4 KB - 64 KB as the keystore size of the logical partition. The value of 0 KB indicates that the keystore function is disabled for the logical partition.

Processor

The **Processor** tab displays the current usage of processors.

Note: When the operating system and the hypervisor support a minimum entitlement of 0.05 processor per virtual processor, the minimum, maximum, and desired processing units can be set to the lowest supported value of 0.05.

Memory

The **Memory** tab displays properties of the running logical partition that is using the dedicated or the shared memory.

Persistent Memory

The **Persistent Memory** tab displays the list of persistent memory volumes that are used by the logical partition. You can also add or remove the persistent memory volume for the logical partition. The Hybrid Memory Subsystem (HMS) view displays information about the persistent memory volumes that are created through the HMS and are associated with the logical partition. You can **Modify**, **Remove**, and **Format** the persistent memory volumes. You can also create a persistent memory volume by using the **Add** menu option under HMS view, which creates a new persistent memory volume and associates the new persistent memory volume with the logical partition.

Note: The Persistent Memory function is currently supported only on the SUSE Linux® operating system.

Physical I/O adapter

The **Physical I/O Adapter** tab displays the properties of all the physical I/O adapters that are available for the managed system and that can be assigned to a partition. You can also add and remove an adapter in a partition.

Change Default Profile

Change the default profile for the partition.

Select a profile from the drop down list to be the new default profile.

Operations

Operations contains the tasks for operating partitions.

About this task

To open the operations tasks that are available for your partitions, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to manage operations tasks. Click **Actions** > **View Partition Properties**
3. In the menu pod, expand **Partition Actions** and then expand **Operations**.
4. Select the operations task that you want to perform from the list.

Activate

Use the **Activate** task to activate a partition on your managed system that is in the **Not Activated** state.

Select the partition profile from the list of profiles and click **OK** to activate the partition. On the **Advanced** tab, select the **No VSI Profile** check box to ignore the failure while configuring the Virtual Station Interface (VSI) profile.

Note: As of HMC Version 7.7, or later, you can install a Virtual I/O Server (VIOS) on a logical partition from an HMC by using a DVD, a saved image, or a Network Installation Management (NIM) server.

Netboot

Use the **netboot** task to network boot an AIX, Linux, or an IBM i partition on your managed system that is in the **Not Activated** state.

The **Network boot** wizard guides you through the steps of installing the operating system on the partition and then activating the partition. Select a partition profile to install the operating system on the partition. Click **Next** to configure the network settings for the logical partition.

Note: For Virtual I/O Server, you must choose the **Install** option from the **Actions** menu to install the VIOS on your managed system that is in the **Not Activated** state.

Restart

Restart the selected logical partition or partitions.

For IBM i logical partitions, use this window only if you cannot restart the IBM i logical partition from the command line of the operating system. Using this window to restart an IBM i logical partition results in an abnormal IPL.

If you choose to restart VIOS partitions that are acting as the Paging Service Partition (PSP) for a number of client partitions, a warning displays, indicating that you must shut down the client partitions before you shut down the VIOS partition.

Choose one of the following options. The Operating System option and the Operating System Immediate option are enabled only if Resource Monitoring and Control (RMC) is up and configured.

Dump
The HMC shuts down the logical partition and initiates a main storage or system memory dump. For AIX and Linux logical partitions, the HMC also notifies the logical partition to shut down. For IBM i logical partitions, the processors are stopped immediately. After the shutdown is complete, the logical partition is immediately restarted. (IBM i logical partitions are restarted multiple times so that the logical partition can store the dump information.) Use this option if a portion of the operation system appears hung and you want a dump of the logical partition for analysis.

Operating System
The HMC shuts down the logical partition normally by issuing a shutdown -r command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. After the shutdown is complete, the logical partition is immediately restarted. This option is only available for AIX logical partitions. Immediate: The HMC shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs that are running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data is partially updated. Use this option only after a controlled end is unsuccessfully attempted.

Operating System Immediate
The HMC shuts down the logical partition immediately by issuing a shutdown -Fr command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. After the shutdown is complete, the logical partition is immediately restarted. This option is only available for AIX logical partitions.

Dump Retry
The HMC retries a main storage or system memory dump on the logical partition. After this operation is complete, the logical partition is shut down and restarted. Use this option only if you previously tried the **Dump** option without success. This option is only available for IBM i logical partitions.

Shut Down

Shut down the selected logical partition or partitions.

For IBM i logical partitions, use this window only if you cannot shut down the IBM i logical partition from the command line of the operating system. Using this window to shut down an IBM i logical partition will result in an abnormal IPL.

If you choose to shut down VIOS partitions that are acting as the Paging Service Partition (PSP) for a number of client partitions, a warning displays, indicating that you should shut down the client partitions before shutting down the VIOS partition.

Choose from the following options:

Delayed
The HMC shuts down the logical partition using the delayed power-off sequence. This allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down

within the predetermined amount of time, it will end abnormally and the next restart may be longer than normal.

Immediate
The HMC shuts down the logical partition immediately. The HMC ends all active jobs immediately. The programs running in those jobs are not allowed to perform any job cleanup. This option might cause undesirable results if data has been partially updated. Use this option only after a controlled shutdown has been unsuccessfully attempted.

Operating System
The HMC shuts down the logical partition normally by issuing a shutdown command to the logical partition. During this operation, the logical partition performs any necessary shutdown activities. This option is only available for AIX logical partitions.

Operating System Immediate
The HMC shuts down the logical partition immediately by issuing a shutdown -F command to the logical partition. During this operation, the logical partition bypasses messages to other users and other shutdown activities. This option is only available for AIX logical partitions.

Delete

Use the **Delete** task to delete the selected partition.

The Delete task deletes the selected partition and all of the partition profiles associated with the partition from the managed system. When you delete a partition, all hardware resources currently assigned to that partition become available to other partitions.

Schedule Operations

Create a schedule for certain operations to be performed on the logical partition.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule an operation to remove resources from a logical partition or move resources from one logical partition to another.

The Scheduled Operations task displays the following information for each operation:

- The processor that is the object of the operation
- The scheduled date
- The scheduled time
- The operation
- The number of remaining repetitions

From the **Scheduled Operations** window you can perform the following operations:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

You can schedule an operation to occur once or you can schedule it to repeat. You must provide the time and date that you want the operation to occur. If you want the operation to repeat, you are asked to select the following time intervals:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)

- The total number of repetitions. (required)

The operations that you can schedule for a logical partition include the following operations:

Activate on an LPAR
Schedules an operation on a selected profile for activation of the selected logical partition.

Dynamic Reconfiguration
Schedules an operation for adding, removing, or moving a resource (processors or megabytes of memory).

Operating System Shutdown (on a partition)
Schedules a shutdown of the selected logical partition.

Backup I/O Configuration
Schedules a backup I/O configuration operation for a selected Virtual I/O Server.

To schedule operations on the HMC, complete the following steps:

1. In the Navigation area, click **Systems Management**.
2. In the work pane, select one or more partitions.
3. In the taskpad, select the **Operations** task category, then click **Schedule Operations**. The **Customize Scheduled Operations** window opens.
4. From the **Customize Scheduled Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, click **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
5. To return to the HMC workplace, point to **Operations** and then click **Exit**.

Validate Maintenance Readiness and Prepare

Use the **Validate Maintenance Readiness** task to validate the readiness of the Virtual I/O Server (VIOS) for maintenance. The VIOS must be in **Running** state with an active Resource Monitoring Control (RMC) connection to perform the validation operation on the VIOS. To complete the validation operation, you must have access to all the partitions of the managed system.

The Hardware Management Console (HMC) validates the readiness of the VIOS for the maintenance. When you execute the maintenance readiness operation, the HMC validates all the client logical partitions that use Virtual I/O Servers for Multi-path I/O operation or redundancy setup for the network and storage that is attached to a logical partition. To check the redundancy setup of the network or storage, the HMC gets the inventory information of other Virtual I/O Servers that are associated with the managed system. However, if other VIOS partitions in the system do not have a proper RMC connection, the validation process continues, and results are shown based on the current states of the Virtual I/O Servers.

The page also displays information about all the impacted client partitions that do not have a redundant Virtual SCSI Storage, Virtual Fibre Channel, Virtual networks, and Virtual NIC that is provided by the VIOS.

You can click **Prepare for Maintenance** in the upper-right corner of the **Validate Maintenance Readiness** window to prepare the VIOS for maintenance. You can select the **Continue with prepare even if there are errors/warnings** checkbox to prepare the VIOS for maintenance even when there are errors and warning.

The **Prepare for Maintenance** task unconfigures the virtual SCSI and virtual Fibre Channel devices in the VIOS to switch the path of the redundant virtual SCSI and the virtual Fibre Channel storage. It also changes the high availability mode of the redundant Shared Ethernet Adapter (SEA) of the VIOS to the Standby state. This task also performs the failover of the vNIC by activating the vNIC backing device that is on an another VIOS. Any failures during these steps are reported at the end of the procedure.

You can click **View System VIOS** in the upper-right corner of the **Validate Maintenance Readiness** window to view information about the Virtual I/O Server of the managed system.

Mobility

Use the Mobility task to migrate your partition to another server, ensure that the requirements for the migration are met, and recover if the partition is in an invalid state.

Migrate

Migrate a partition to another managed system.

About this task

To migrate a partition to another system, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. In the content pane, select the server. Click **Actions** > **View System Partitions**.
3. In the content pane, select the partition that you want to migrate to another system.
4. Click **Actions** > **Mobility** > **Migrate**. The Partition Migration wizard opens.
5. Complete the steps in the Partition Migration wizard and click **Finish**.

Validate

Validate the settings for moving the partition from the source system to the destination system.

About this task

To validate the settings, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. In the content pane, select the server. Click **Actions** > **View System Partitions**.
3. In the content pane, select the partition for which you want to validate the settings.
4. Click **Actions** > **Mobility** > **Validate**. The Partition Migration Validation window opens.
5. Fill in the information in the fields, and click **Validate**.

Recover

Recover this partition from a migration that did not complete.

About this task

To recover this partition from a migration that did not complete, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. In the content pane, select the server. Click **Actions** > **View System Partitions**.
3. In the content pane, select the partition that you want to recover.
4. Click **Actions** > **Mobility** > **Recover**. The Migration Recovery window opens.

5. Complete the information as necessary and click **Recover**.

Partition Templates

Partition templates contain details for partition resources, such as physical adapters, virtual networks, and storage configuration. You can create client partitions from the quick-start templates that are available in the template library or from your own user-defined templates on the Hardware Management Console (HMC).

Capture Partition as a Template

You can capture the configuration details of a running partition and save the information as a partition template by using the Hardware Management Console (HMC).

To capture the configuration as a template, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to capture as a template.
3. Click **Actions** > **Templates** > **Capture Partition as a Template**.
4. Enter a template name and description.
5. Click **OK**.

Use the online Help if you need additional information about this task.

Profiles

Learn about the tasks that are available in the **Profiles** menu.

Manage Profiles

Use the **Manage Profiles** task to create, edit, copy, delete, or activate a profile for the selected partition.

A partition profile contains the resource configuration for the partition. You can modify the processor, memory, and adapter assignments for a profile by editing the profile.

The default partition profile for a logical partition is the partition profile that is used to activate the logical partition if no other partition profile is selected. You cannot delete the default partition profile unless you first designate another partition profile as the default partition profile. The default profile is defined in the status column.

Choose **Copy** to create an exact copy of the selected partition profile. This allows you to create multiple partition profiles that are nearly identical to one another by copying a partition profile and changing the copies as needed.

Manage Custom Groups

Groups are comprised of logical collections of objects. You can report status on a group basis, allowing you to monitor your system in a way that you prefer. You can also nest groups (a group contained within a group) to provide hierarchical or topology views.

One or more user-defined groups might already be defined on your Hardware Management Console (HMC). Default groups are listed under **Custom Groups** node under **Configuration**. The default groups are **All Partitions** and **All Objects**. You can create others, delete the ones that were created, add to created groups, create groups using the pattern match method, or delete from created groups by using the **Manage Custom Groups** task.

Use the online Help if you need additional information for managing custom groups.

Save Current Configuration

Save the current configuration of a logical partition to a new partition profile by entering a new profile name.

This procedure is useful if you change the configuration of a logical partition using dynamic logical partitioning and you do not want to lose the changes when you restart the logical partition. You can perform this procedure at any time after you initially activate a logical partition.

Delete Partition

You can delete a partition and the associated partition profile by using the Hardware Management Console (HMC).

To delete a partition, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to delete.
3. Click **Actions** > **Delete Partition**.
4. Select any options that you want.
5. Click **OK**.

Use the online Help if you need additional information about this task.

Serviceability

Problem Analysis on the HMC automatically detects error conditions and reports to you any problem that requires service to repair it.

These problems are reported to you as serviceable events. Use the **Serviceable Events Manager** task to view specific events for selected systems. However, if you notice a problem occurred or you suspect a problem is affecting the system but Problem Analysis has not reported it to you, use the **Create Serviceable Event** task to report the problem to your service provider.

Serviceable Events Manager

Problems on your managed partitions are reported to the Hardware Management Console (HMC) as serviceable events. You can view the problem, manage problem data, call home the event to your service provider, or repair the problem.

To set the criteria for the serviceable events you to view, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to manage serviceable events.
3. In the menu pod, expand **Serviceability** and then click **Serviceability**.
4. Click **Serviceable Events Manager**.
5. Provide event criteria, error criteria, and FRU criteria. If you do not want the results that are filtered, select **ALL**. Click **Refresh** to refresh the list of serviceable events based on the criteria filter values.

The **Serviceable Events Overview** window displays all of the events that match your criteria. The information that is displayed in the compact table view includes the following information:

- Problem Number
- PMH Number
- Reference Code - Click the Reference code to display a description of the problem reported and actions that can be taken to fix the problem.
- Status of the problem
- Last reported time of the problem
- Failing MTMS of the problem

- Originating HMC

The full table view includes more detailed information, including reporting partition ID, primary data event timestamp, duplicate count, notification type, first reported time, reporting name, reporting MTMS, firmware fix, and serviceable event text.

Select a serviceable event and use the **Action** drop-down menu to:

- **View Details**: Field-replaceable units (FRUs) associated with this event and their descriptions.
- **View Files**: View the files associated with the selected serviceable event.
- **View Reference Code Description**: View the description of the reference code associated with the selected serviceable event. The option will not be available if additional description is not available.
- **Call Home**: Report the event to your service provider.
- **Repair**: Start a guided repair procedure, if available.
- **Close Event**: After the problem is solved, add comments and close the event.
- **Add PMH Comment**: Add a comment to a PMH for a selected serviceable event. If a PMH number does not exist for a given problem, the Add PMH Comment option is not available.

Use the online Help if you need additional information on managing serviceable events.

Reference Code Log

Use the **Reference Code Log** task to view reference codes that have been generated for the selected logical partition. Reference codes are diagnostic aids that help you determine the source of a hardware or operating system problem.

By default, only the most recent reference codes that the logical partition has generated are displayed. To view more reference codes, enter the number of reference codes that you want to view into **View history** and click **Refresh**. The window displays that number of the latest reference codes, with the date and time at which each reference code was generated. The window can display up to the maximum number of reference codes stored for the logical partition.

Control Panel Functions

This task displays the available virtual control panel functions for the selected IBM i partition. The tasks are:

(21) Activate Dedicated Service Tools
Starts Dedicated Service Tools (DST) on the partition.

(65) Disable Remote Service
Deactivates remote service on the partition.

(66) Enable Remote Service
Activates remote service on the partition.

(68) Concurrent Maintenance Power Off Domain
Concurrent maintenance power domain Power Off.

(69) Concurrent Maintenance Power On Domain
Concurrent maintenance power domain Power On.

Virtual I/O

Learn how to view the virtual networks, virtual network interface controllers, and virtual storage of a partition.

You can use the Hardware Management Console (HMC) to view the virtual topology of a partition.

Virtual Networks

You can view and add virtual networks that are associated with the selected logical partition.

The **Virtual Networks** table lists the virtual network name, VLAN ID, virtual switch, virtual network bridge, and virtual Ethernet adapter ID that are associated with each virtual network. You can click **Attach Virtual Network** to view the available virtual networks and attach additional virtual networks to the logical partition.

To view the virtual networks for the selected partitions by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Virtual Networks**.

Use the online Help if you need additional information about this task.

Virtual NIC

You can manage all aspects of the virtual Network Interface Controller (NIC) configuration that is associated with the partition.

A virtual NIC is a type of virtual adapter that can be configured on logical partitions to provide a network interface. Each virtual NIC client adapter is backed by an SR-IOV logical port that is owned by the hosting partition.

The **Virtual NIC** table lists all virtual NICs that are configured for the selected partition. A virtual NIC can have one or more backing devices. The maximum number of backing devices per virtual NIC depends on the system. If the virtual NIC has more than one backing device, you can expand the node to view all the backing devices. If the virtual NIC has only one backing device, that backing device is the active backing device. The active backing device is the one that is in use by the virtual NIC. If the managed system is not failover capable, the table displays virtual NICs that have a single backing device.

You can add a virtual NIC to the partition. To add a virtual NIC, click **Add Virtual NIC**. You can configure the virtual NIC only in dedicated mode. You can also modify and view virtual NIC properties. To modify properties of a virtual NIC, select the virtual NIC in the table and click **Action** > **Modify vNIC** . To view the properties of a virtual NIC, select the virtual NIC in the table and click **Action** > **View vNIC**.

To view the virtual NIC for the selected partition by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Virtual NICs**.

Use the online Help if you need additional information about this task.

Virtual Storage

You can create, view, and manage the storage capability of the logical partition.

The **Virtual Storage** table displays the Virtual Small Computer Serial Interface (SCSI) devices that are configured to a logical partition. You can also view the information about the physical volume groups, shared storage pool volume, and the logical volume.

You can add the virtual storage resources to a partition. Click **Adapter View** to create, view the adapter configuration of the virtual storage devices that are allocated for the logical partition. Click **Storage View** to view and manage the storage capability of the logical partition.

Physical volumes can be exported to partitions as virtual SCSI disks. Click **Show assigned physical volumes** to view the assigned physical volumes that are assigned to the logical partition.

To add physical volumes to a partition, select the physical volumes from the list and specify the **User Defined Name** for each physical volume that you want to add to the partition and then click **OK**. If you want to change the server adapter ID that is assigned to each physical volume, click **Edit** for each of the physical volumes that you want to update. The **Edit connection** window is displayed. You can specify up to 3 Virtual I/O servers, and then enter the new server adapter ID that you want to assign for the adapter connection.

To add different types of virtual storage devices to a partition, click **Add Virtual SCSI Device**. Select the available virtual storage that you want to add. You can select the virtual storage types such as **Physical Volume**, **Shared Storage Pool Volume**, or **Logical Volume**.

To view the virtual storage for the selected partition by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Virtual Storage**.

Use the online Help if you need additional information about this task.

Hardware Virtualized I/O

You can view and change the settings of hardware virtualized I/O adapters, such as single root I/O virtualization (SR-IOV) port adapters for a partition by using the Hardware Management Console (HMC).

To view the hardware virtualized I/O adapters for the selected partition by using the HMC, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Partitions**.
2. Select the partition for which you want to manage serviceability tasks and click **Actions** > **View Partition Properties**.
3. In the menu pod, expand **Virtual I/O** and then click **Hardware Virtualized I/O**.
4. You can add an SR-IOV logical port to the partition or change the settings of the SR-IOV logical ports. In the **SR-IOV logical port** table, you can also view the information about the logical ports that can be migrated and the information about the backing device that is configured for the logical ports.

   Notes:

   - With HMC Version 9.1.930, or later, the HMC also supports the RDMA over Converged Ethernet (RoCE) adapter.
   - If you are using HMC Version 9.1.940, with firmware at level FW940, or later, you can create logical partitions that have an SR-IOV logical port that can be migrated. You can migrate a logical partition with SR-IOV logical ports when the Migratable option is used to create a backup virtual device when creating a logical port. The backup device can be either a virtual Ethernet or a virtual Network Interface Controller (NIC) adapter. When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

   Use the online Help if you need additional information about this task.

Manage Groups

The **All groups** view provides a mechanism for you to group system resources together in a single view.

Groups may be nested to create customized system resources view.

You can view all the groups that are created by the users of the management console, including cumulative state information for system resources in a group. A custom group can consist of any systems, partitions, and Virtual I/O Servers that are managed by the management console.

To create a new group, complete the following steps:

1. Click **Create group** on the toolbar.
2. In the **Create group** window, specify a group name and description for the group. You can also tag a color to the group that you want to create.
3. Select one or more resources (for example: servers, partitions, or frames) that you want to include in the group that you want to work with.
4. Click **OK** to save the changes and to close the window.

You can edit an existing group to add or remove the resources from the group.

Note: When the last member (resources) of the group is removed, a message is displayed to confirm whether you want to delete the group. Click **Cancel** to retain the group in the **All groups** view.

Power Enterprise Pools

Systems Management for Power Enterprise Pool displays Power Enterprise Pool tasks that you can perform.

You can perform the following operations by using the Power Enterprise Pool offering:

- Add processors or memory to a server.
- Remove processors or memory from a server.
- Update the pool configuration.
- Add a server to the pool.
- Remove an existing server from the pool.
- Add processors or memory to the pool.
- View the following Power Enterprise Pool information:
  - Pool membership information
  - Pool resource information
  - Pool compliance information
  - Pool history log

Console Management tasks

Learn about the tasks that are available on the Hardware Management Console (HMC) under **Console Management**.

To open these tasks, see "HMC tasks, user roles, IDs, and associated commands" on page 8.

Note: Depending on the task roles that are assigned to your user ID, you might not have access to all the tasks. See Table 3 on page 8 for a listing of the tasks and the user roles that are allowed to access them.

Launch Guided Setup Wizard

This task uses a wizard to set up your system and HMC.

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Launch Guided Setup Wizard**.
3. From the **Launch Guided Setup Wizard - Welcome** window it is recommended that you have certain prerequisites on hand. Click **Prerequisites** in the **Launch Guided Setup Wizard - Welcome** window for the information. When you have completed that, this wizard takes you through the following tasks required to set up your system and HMC. As you complete each task, click **Next** to proceed.

   a. Change HMC Date and Time
   b. Change HMC passwords
   c. Create additional HMC users
   d. Configure HMC Network Settings (This task cannot be performed if you are accessing the **Launch Guided Setup Wizard** remotely.)
   e. Specify contact information
   f. Configure connectivity information
   g. Authorize users to use the Electronic Service Agent software tool and configure notification of problem events.
4. Click **Finish** when you have completed all the tasks in the wizard.

View Network Topology

This task allows you to view and ping the connectivity between various network nodes within the Hardware Management Console (HMC).

To view the network topology, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **View Network Topology**.
3. From the **View Network Topology** window, you can ping current and saved nodes.
4. Click **Close** when you have completed this task.

Use the online Help if you need additional information about viewing the network topology.

Test Network Connectivity

This task allows you to view network diagnostic information about the network protocols for the Hardware Management Console (HMC).

To test the network connectivity, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Test Network Connectivity**.
3. From the **Test Network Connectivity** window, you can work with the following tabs:

   Ping
   You can ping the TCP/IP address or name.

   Interfaces
   Displays the statistics for the network interfaces that are currently configured. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Ethernet Settings
   Displays the settings for the Ethernet cards that are currently configured. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Address
   Display the TCP/IP addresses for the configured network interfaces. To update the information that is currently displayed with the most recent information, click **Refresh**.

   Routes
   Displays the Kernel IP and IPv6 routing tables and corresponding network interfaces. To update the information that is currently displayed with the most recent information, click **Refresh**.

ARP
Displays the contents of the Address Resolution Protocol (ARP) connections. To update the information that is currently displayed with the most recent information, click **Refresh**.

Sockets
Displays information about TCP/IP sockets. To update the information that is currently displayed with the most recent information, click **Refresh**.

TCP
Displays information about Transmission Control Protocol (TCP) connections. To update the information that is currently displayed with the most recent information, click **Refresh**.

IP Tables
Displays information (in table format) about the Internet Protocol (IP) packet filter rules. To update the information that is currently displayed with the most recent information, click **Refresh**.

UDP
Displays information about User Datagram Protocol (UDP) statistics. To update the information that is currently displayed with the most recent information, click **Refresh**.

4. Click **Cancel** when you have completed this task.

Use the online Help if you need additional information about testing the network connectivity.

Change Network Settings

This task allows you to view the current network information for the HMC and to change network settings.

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change Network Settings**.
3. From the **Change Network Settings** window, you can work with the following tabs:

Identification
Contains the host name, domain name, and console description of the HMC.

LAN Adapters
A summarized list of all (visible) Local Area Network (LAN) adapters. You can select any of these and click **Details...** to open a window allowing you to change addressing, routing, other LAN adapter characteristics, and firewall settings.

Bond LAN Adapters
Create or delete a Bond LAN adapter. A Bond LAN adapter combines two Ethernet interfaces into a single logical link. To change the settings of the Bond LAN adapter, select a Bond LAN adapter and click **Edit**. You can change the IP address, IP network mask, gateway, and the firewall settings of the Bond LAN adapter.

Name Services
Specify the DNS and domain suffix values for configuring the console network settings.

Routing
Specify the routing information and default gateway information for configuring the console network settings.

The **Gateway address** is the route to all networks. The default gateway address (if defined) informs this HMC where to send data if the target station does not reside on the same subnet as the source. If your machine can reach all stations on the same subnet (usually a building or a sector within a building), but cannot communicate outside the area, it is usually because of an incorrectly configured default gateway.

You can assign a specific LAN to be the **Gateway device** or you can choose “any.”

You can select **Enable ’routed’** to start the routed daemon, which allows it to run and allows any routing information to be exported from the HMC.

4. Click **OK** when you have completed this task.

Note: Depending on the type of change that you make, the network or console automatically restarts or the console automatically reboots.

Use the online Help to get additional information for customizing the network settings.

Change Performance Monitoring Settings

The Performance and Capacity Monitor tool collects allocation and usage data for virtualized server resources. It displays data in the form of graphs and tables, which are viewable from the Performance and Capacity Monitor home page.

The Performance and Capacity Monitor gathers data and provides capacity reporting and performance monitoring. This information can help you to determine the available capacity and whether your resources might be overextended or underused. In addition, your interpretation of the graphs and tables might be useful for capacity planning and troubleshooting. For more information about The Performance and Capacity Monitor tool, see Using the Performance and Capacity Monitor.

The Performance and Capacity Monitor captures data only from the servers for which you choose to enable data collection.

To enable data collection, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change Performance Monitoring Settings**.
3. Specify the number of days for which you want to store performance data by typing in a number 1 - 366. Alternatively, you can click the up or down arrows next to **Number of days to store performance data** under **Performance Data Storage**.

   Note: By default, the HMC stores data for 180 days. However, you can specify the maximum number of days that the HMC stores data to 366 days.
4. Click the toggle switch in the **Collection** column next to the name of the server for which you want to collect data. Alternatively, you can click **All On** to enable data collection for all of the servers in your environment that the HMC manages.

   Note: You might be prevented from collecting data from all of the servers in your environment because storage space is limited. The HMC prohibits you from enabling data collection from more servers when the HMC determines that it might run out of estimated storage space.
5. Click **OK** to apply the changes and close the window. You can now review the collected data when you access the Performance and Capacity Monitor home page.

Change Date and Time

Change the time and date of the battery-operated Hardware Management Console (HMC) clock and add or remove time servers for the Network Time Protocol (NTP) service.

Use this task in the following situations:

- If the battery is replaced in the HMC.
- If your system is physically moved to a different time zone.

Note: The time setting adjusts automatically for Daylight Saving Time in the time zone you select.

To change the date and time, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change Date and Time**.
3. Click the **Customize Console Date and Time** tab.
4. Enter the date and time information.
5. Click **OK**.

To change the time server information, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change Date and Time**.
3. Click the **NTP Configuration** tab.
4. Provide the appropriate information for the time server.
5. Click **OK**.

If you need additional information for changing the date and time of the HMC or for adding or removing time servers for the Network Time Protocol (NTP) service, use the online Help.

Change Language and Locale

This task sets the language and location for the HMC. After you select a language, you can select a locale associated with that language.

The language and locale settings determine the language, the character set, and other settings specific to the country or region (such as formats for date, time, numbers, and monetary units). Changes made in the **Change Language and Locale** window affect only the language and locale for the HMC itself. If you access the HMC remotely, the language and locale settings on your browser determine the settings that the browser uses to display the HMC interface.

To change the language and locale on the HMC:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Change Language and Locale**.
3. From the **Change Language and Locale** window, choose the applicable language and locale.
4. Click **OK** to apply the change.

Use the online Help if you need additional information for changing the language and locale of the HMC.

Create Welcome Text

Create and display a welcome message or display a warning message that appears before users log on to the Hardware Management Console (HMC).

The text that you enter in the message input area for this task appears on the **Welcome** window after you initially access the console. You can use this text to notify users about certain corporate policies or security restrictions that apply to the system.

To create a welcome text, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.
2. In the content pane, click **Create Welcome Text**.
3. Enter the welcome text that you want to display in the text box.

   Note: A maximum of 8192 characters is allowed.
4. Click **OK**.

For more information about this task, use the online Help.

Console Default Settings

You can modify the default console settings on the Hardware Management Console (HMC).

You can also modify the number of days for which a certificate is valid.

Note: The certificate can be valid for maximum of 3650 days.

To modify the console default settings, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Settings**.

2. In the content pane, click **Console Default Settings**.
3. In the **Console Default Settings** window, you can specify the number of days for which the certificate is valid. You can also configure the following user settings:
   - `Idle Time Out` to specify idle time out settings for an HMC session in minutes.
   - `Session Time out` to specify session time out settings for an HMC session in minutes.
   - `Max Web UI Login Attempts` to specify the maximum number of login attempts to the HMC graphical user interface (GUI). You can enter a value in the range of 3 - 50.
   - `Web UI Suspend Time` to specify the login suspension time for HMC GUI. You can enter a value in the range of 1 - 1440 minutes

     Note: After the maximum number of failed login attempts is reached, the account is suspended for the number of specified minutes in the `Web UI Suspend Time` field.
4. When you complete the task, click **OK**.

Use the online Help if you need additional information about this task.

Shut Down or Restart

This task enables you to shut down (power off the console) or to restart the console.

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Shut Down or Restart**.
3. From the **Shut Down or Restart** window, you can:
   - Select **Restart the HMC** to automatically restart the HMC once the shut down has occurred.
   - Do not select **Restart the HMC** if you do not want to automatically restart the HMC.
4. Click **OK** to proceed with the shut down, otherwise click **Cancel** to exit the task.

Use the online Help if you need additional information about shutting down or restarting the HMC.

Schedule Operations

Create a schedule for certain operations to be performed on the Hardware Management Console (HMC) itself without operator assistance.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule a backup of important HMC information to DVD to occur once, or set up a repeating schedule.

The **Scheduled Operations** task displays the following information for each operation:

- The processor that is the object of the operation.
- The scheduled date.
- The scheduled time.
- The operation.
- The number of remaining repetitions.

From the **Scheduled Operations** window you can:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.

- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

An operation can be scheduled to occur one time or it can be scheduled to be repeated. You are required to provide the time and date that you want the operation to occur. If the operation is scheduled to be repeated, you are asked to select:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)
- The total number of repetitions. (required)

The operation that can be scheduled for the HMC is:

Backup Critical Console Data
Schedules an operation to back up the critical console hard disk information for the HMC.

To schedule operations on the HMC, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Schedule Operations**.
3. From the **Schedule Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, point to **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
4. To return to the HMC workplace, point to **Options** and then click **Exit**.

Use the online Help to get additional information for scheduling an operation.

View Licenses

View the Licensed Internal Code that you agreed to for this Hardware Management Console (HMC).

You can view licenses at any time. To view licenses, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **View Licenses**.
3. Click any of the license links to view more information.

   Note: This list does not include programs and code that is provided under separate license agreements.
4. Click **OK**.

Update the Hardware Management Console

Learn how to update the internal code of the Hardware Management Console (HMC) and view system information and system readiness.

To update the HMC, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Update the Hardware Management Console**. The **Install HMC Corrective Service Wizard** opens.
3. Click **Next** to start the update process.
4. Follow the steps in the wizard to complete the update operation.
5. Click **Finish** when you have completed this task.

Use the online Help if you need additional information about updating the Hardware Management Console.

Format Media

This task formats a diskette or USB 2.0 Flash Drive Memory Key.

You can format a diskette by supplying a user-specified label.

To format a diskette or USB 2.0 Flash Drive Memory Key, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Format Media**.
3. From the **Format Media** window, select the type of media you want to format, then click **OK**.
4. Make sure that your media is correctly inserted, then click **Format**. The **Format Media** progress window is displayed. When the media is formatted, the **Format Media Completed** window is displayed.
5. Click **OK** and then click **Close** to end the task.

Use the online Help if you need additional information about this task.

Backup Management Console Data

This task backs up (or archives) the data that is stored on your Hardware Management Console (HMC) hard disk that is critical to support HMC operations.

Back up the HMC data after changes are made to the HMC or information that is associated with logical partitions.

The HMC data that is stored on the HMC hard disk drive can be saved to a DVD-RAM on a local system, a remote system that is mounted to the HMC file system (such as NFS), or sent to a remote site by using File Transfer Protocol (FTP).

By using the HMC, you can back up all important data, such as the following data:

- User-preference files
- User information
- HMC platform-configuration files
- HMC log files
- HMC updates through Install Corrective Service.

Note: Use the archived data only along with a reinstallation of the HMC from the product CDs.

To back up the HMC critical data, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Backup Management Console Data**.
3. From the **Backup Management Console Data** window, choose the archive option that you want to complete.

4. Click **Next**, then follow the appropriate instructions that are associated with the option you chose.
5. Click **OK** to continue with the backup process.

Use the online Help if you need additional information for backing up the HMC data.

Note:

- <li>For the HMC model 7063-CR1, DVD media is not supported.
- If you are using HMC Version 9.1.940, or later, you can specify a name for the generated backup file. If the backup file exists on the server, select the **Replace file** to replace the contents of the existing file that has the same name.

Restore Management Console Data

This task is used to select a remote repository for restoring critical backup data for the HMC.

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Restore Management Console Data**.
3. From the **Restore Management Console Data** window, click **Restore from a remote Network File System (NFS) server**, **Restore from a remote File Transfer Protocol (FTP) server**, **Restore from a remote Secure Shell File Transfer Protocol (SFTP) server**, or **Restore from a remote removable media**.
4. Click **Next** to proceed or **Cancel** to exit the task without making any changes.

Use the online Help if you need additional information about restoring critical backup data for this HMC.

Save Upgrade Data

This task uses a wizard to save upgrade data to selected media. This data consists of files that were created or customized while running the current software level. Saving this data to selected media is performed prior to an HMC software upgrade.

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Save Upgrade Data**.
3. From the **Save Upgrade Data** window, this wizard takes you through the steps required for saving your data. Select the type of media you want to save your data to, then click **Next** to proceed through the task windows.
4. Click **Finish** when you have completed the task.

Use the online Help if you need additional information for saving upgrade data.

Manage Data Replication

This task enables or disables customized data replication. Customized data replication allows another HMC to obtain customized console data from or send data to this HMC.

The following types of data can be configured:

- User profiles data
  - User identification
  - Authentication methods
  - Resource roles and task roles that are managed by the user
  - Logon session properties
  - Remote access settings

- Group data
  - All user-defined group definitions

  Note: When you configure **Group Data**, the complete group data is transferred from the source HMC and replaced on the secondary HMC. If the secondary HMC does not have a resource that is part of a group, the resource is not shown in that particular group.
- Multi-Factor Authentication data
  - PowerSC MFA hostname that is used by the HMC for Multi-Factor Authentication.
- Kerberos configuration data
  - Key Distribution Center (KDC), realm, and hostname that is used by the HMC for Kerberos authentication
- LDAP configuration data
  - LDAP server name and distinguished name tree that is used by the HMC for LDAP authentication.
- Password policy configuration data
  - Password policy name
  - Password policy description
  - Configured attributes of the password policy (such as min_pwage, pwage, min_length, hist_size, warn_pwage, min_digits, min_uppercase, min_lowercase, and min_special_chars)
- Outbound connectivity data
  - Information for dialing out (such as whether to enable the local system as a call-home server, or whether to allow dialing to use the local modem, the dial prefix, or phone numbers)

Note: Customizable console data is accepted from other HMCs only after specific HMCs and their associated allowable customizable data types are configured.

To manage data replication, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Management**.
2. In the content pane, click **Manage Data Replication**.
3. From the **Manage Data Replication** window, choose the appropriate option that you want to perform.

Use the online Help to get additional information for enabling or disabling customizable data replication.

Templates and OS Images

System templates contain configuration details for resources such as system properties, shared processor pools, reserved storage pool, shared memory pool, Host Ethernet adapters, and SR-IOV adapters. Many of the system settings that you previously configured by using separate tasks are available in the Deploy System from Template wizard. For example, you can configure the Virtual I/O Servers, virtual network bridges, and virtual storage settings when you use the wizard to deploy a system from a system template.

The template library includes predefined system templates, which contain configuration settings based on common usage scenarios. Predefined system templates are available for your immediate use. You can view, modify, deploy, copy, import, export, or delete templates that are available in the template library.

You can also create custom system templates that contain configuration settings that are specific to your environment. You can create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration.

To access the Template Library, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Templates and OS Images**.
2. From the **Templates and OS Images** window, you can access:

- **System**
- **Partition**
- **OS and VIOS Images**

3. When you complete this task, click **Close**.

System Templates

System templates contain configuration information about resources such shared processor pools, reserved storage pool, shared memory pool, physical I/O adapters, Host Ethernet adapters, single root I/O virtualization (SRIOV) adapters, Virtual I/O Server (VIOS), virtual networks, and virtual storage.

You can create custom system templates that contain configuration settings that are specific to your environment. You can also create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration. Click the template name to see the details about the template. Select a system template from the list to view, edit, copy, delete, deploy, or export a template.

Use the online Help if you need additional information on system templates.

Partition Templates

Partition templates contain details about partition resources, such as physical adapters, virtual networks, and storage configuration.

You can create custom partition templates that contain configuration settings that are specific to your environment. You can also create a custom template by copying a predefined template and changing it to fit your needs. Or, you can capture the configuration of an existing system and save the details in a template. Then, you can deploy that template to other systems that require the same configuration. Click the template name to see the details about the template. Select a partition template from the list to view, edit, copy, delete, deploy, or export a template.

Notes:

- If you are using HMC Version 9.1.940, or later, and if you are using a non-captured template to create a logical partition, you can configure an SR-IOV logical port that can be migrated. Select **migratable** in the **Edit** menu of the partition template. You can migrate the logical partition by using the SR-IOV logical port by creating a backup device and associate the SR-IOV logical port to the logical partition. The backup device can either be a virtual Ethernet adapter or a virtual Network Interface Controller (NIC) adapter.
- When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Use the online Help if you need additional information on partition templates.

VIOS Images

Define VIOS images and installation resources for the operating environment that the Hardware Management Console (HMC) can access and use.

You can access the following tasks:

Manage Virtual I/O Server Image Repository

As of HMC version 7.7, or later, you can store the Virtual I/O Server (VIOS) images from a DVD, a saved image, or a Network Installation Management (NIM) server on the HMC. The stored VIOS images can be used for VIOS installation. You must be an HMC super administrator (hmcsuperadmin) to install the VIOS image.

About this task

To manage or to import the VIOS image repository, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Templates and OS Images**.
2. From the **Templates and OS Images** window, select the **OS and VIOS Images** tab, and then click **Manage Virtual I/O Server Image Repository**.
3. In the Virtual I/O Server Image Repository window, click **Import New Virtual I/O Server Image**.
4. In the Import New Virtual I/O Server Image window, choose to import the VIOS images from a DVD or from a file system.

   - To import VIOS images from a DVD to the HMC, complete the following steps:

     a. In the Import Virtual I/O Server Image window, select **Management console DVD**.

     b. In the **Name** field, enter the VIOS image name that you want to import from the DVD.

     c. Click **OK**.

   - To import VIOS images from a Network File System (NFS), File Transfer Protocol (FTP), or Secure Shell File Transfer Protocol (SFTP), complete the following steps:

     a. In the Import Virtual I/O Server Image window, select **File System**.

     b. Select **Remote NFS Server**, **Remote FTP Server**, or **Remote SFTP Server**.

     c. Enter the required details and click **OK**.

Manage Virtual I/O Server Backups

With HMC version 9.2.950, or later, you can manage the I/O configuration of Virtual I/O Servers and manage the backup of the VIOS image on the management console.

About this task

To manage the backup or restore operation of the I/O configuration of the VIOS and to manage the VIOS image, complete the following steps:

Procedure

1. In the navigation area, click **Console Management**, and then select **Templates and OS Images**.
2. From the **Templates and OS Images** window, select the **VIOS Images** tab, and then click **Manage Virtual I/O Server Backups**.
3. In the **Manage Virtual I/O Server Backups** window, select the **Virtual I/O Server Configuration Backup** tab. A table is displayed that lists all the backup files of the VIOS configuration that is taken by the HMC. Additionally, you can view the time at which the configuration file was last edited.

   a) To take the backup of the input/output configuration of a VIOS, click **Backup I/O configuration**. In the Backup I/O configuration window, select the managed system and the VIOS for which the back up is created, and then specify a name for the backup file.

   The name you specify must consist of 1 - 40 characters including file extension **`.tar.gz`**. You can use the characters A - Z and a - z, the numbers of 0 - 9, the dot (.), the dash (-) and the underscore (_) characters.

   b) To rename an existing backup file that is stored in the HMC, select a configuration file from the table and click **Action** > **Rename**.

   c) To restore the VIOS input/output configuration, select a backup file which contains the I/O configuration of the VIOS that you want to restore, and click **Action** > **Restore**.

   d) To export the backup files to the remote system, select one or more Virtual I/O Server backup files that are saved on the HMC and click **Action** > **Export**. You can export the backup files to the remote systems by specifying the proper credentials of the remote server where the backup files are copied.

   e) To import the backup files from the remote system, click **Import**. After specifying the proper credentials of the remote server, select the Virtual I/O Server where the backup files are imported.

4. In the **Manage Virtual I/O Server Backups** window, click the **Virtual I/O Server Backup** tab. A table is displayed that list all the VIOS image backup that are taken in the HMC. Additionally, you can also view the name and size of the VIOS image, the time when the VIOS image file was last edited, the managed system and the VIOS from which the image was captured.

   a) To take the backup of the VIOS image, click **Create Backup**. In the Create Backup window, select the managed system and the VIOS for which the back up is created, and then specify a name for the backup file.

      The name you specify must consist of 1 - 40 characters including file extension `.tar`. You can use the characters A - Z and a - z, the numbers of 0 - 9, the dot (.), the dash (-) and the underscore (_) characters.

   b) To rename an existing VIOS image backup file that is stored in the HMC, select a backup file from the table and click **Action** > **Rename**.

   c) To remove a VIOS image backup file from the HMC, select a backup file which contains the VIOS configuration that you want to remove from the table, and click **Action** > **Remove**.

   d) To export the Virtual I/O Server backup files to the remote system, select one or more backup files that are saved on the HMC and click **Action** > **Export**. You can export the backup files to the remote systems by specifying the proper credentials of the remote server where the backup files are copied.

   e) To import the backup files from the remote system, click **Import**. After specifying the proper credentials of the remote server, select the target Virtual I/O Server where the backup files are imported.

5. In the **Manage Virtual I/O Server Backups** window, click the **Shared Storage Pool Cluster Backup** tab. A table is displayed that lists all the VIOS image backups that are taken in the HMC. Additionally, you can also view the **Managed System, Virtual I/O Server**, and **Backup Created time**.

   a) To take a back up of the VIOS Shared Storage Pool Cluster (SSP) configuration, click **Create Shared Storage Pool Cluster Backup**. In the **Create SSP Cluster Backup** window, select the Managed System and the Virtual I/O Server that the backup is created for, and then specify a name for it.

      The name that you specify must consist of 1 - 40 characters including file extension `.tar.gz`. You can use the characters A - Z and a - z, the numbers of 0 - 9, the dot (.), the dash (-) and the underscore (_) characters.

   b) To restore the Shared Storage Pool Cluster configuration, select the backup file that contains the configuration of the SSP cluster that you want to restore, and click **Action** > **Restore**.

   c) To rename one of the existing SSP cluster configuration backup files, select a configuration from the table and click **Action** > **Rename**.

   d) To remove a SSP cluster configuration backup, select a backup file which contains the SSP cluster configuration that you want to remove, and click **Action > Delete**.

   e) To export the Virtual I/O Server backup files to the remote system, select one or more backup files that are saved on the HMC and click **Action** > **Export**. You can export the backup files to the remote systems by specifying the proper credentials of the remote server where the backup files are copied.

   f) To import the backup files from the remote system, click **Import**. After specifying the proper credentials of the remote server, select the target Virtual I/O Server where the backup files are imported.

6. Click **OK**.

All System Plans

A system plan is a specification of the logical partition configuration of a single managed system.

The table lists all the system plans that can be used to configure a managed system. You can create your own system plan or import an existing system plan.

Create System Plan

You can create a new system plan for a system that this Hardware Management Console (HMC) manages. The new system plan contains specifications for the logical partitions and partition profiles of the managed system that you used to create the plan.

1. Click **Create**.
2. Select a managed system from the available list and complete the **System plan name** and **Plan description** fields.
3. Check any options that you want.
4. Click **Create**.

Import System Plan

You can import a system plan file to the Hardware Management Console (HMC). The new system plan contains specifications for the logical partitions and partition profiles of the managed system that you used to create the plan.

1. Click **Import**.
2. Select a source to import the system plan file to the HMC.
3. Click **Import**.

Export System Plan

You can export a system plan file from the Hardware Management Console (HMC).

1. Select the system plan from the list and click **Actions** > **Export**.
2. Select a source to export the system plan file to the HMC.
3. Click **Export**.

Deploy System Plan

You can deploy a system plan file to one or more systems that the HMC manages. The managed system that you deploy the system plan on must have hardware that is identical to the hardware in the system plan.

1. Select the system plan from the list and click **Actions** > **Deploy**.
2. Follow the instructions on the **Deploy System Plan** wizard.

Delete System Plan

You can delete a system plan file from the Hardware Management Console (HMC).

1. Select the system plan from the list and click **Actions** > **Delete**.

Refresh

You can refresh the table to see any recent changes to the available system plans.

1. Click **Refresh** to update the table with the latest data.

Use the online Help if you need additional information about this task.

Users and Security tasks

The tasks that are available on the HMC for the **Users and Security** tasks are described.

Note: Depending on the task roles assigned to your user ID you may not have access to all the tasks. See "HMC tasks, user roles, IDs, and associated commands" on page 8 for a listing of the tasks and the user roles allowed to access them.

Change User Password

This task allows you to change your existing password that is used for logging on to the Hardware Management Console (HMC). A password verifies your user ID and your authority to log in to the console.

To change your password, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Change User Password**.
3. From the **Change User Password** window, specify your current password, specify a new password that you want to use, and re-specify the new password for confirmation in the fields provided.

   Note: The new password that you specify must have at least eight characters.
4. Click **OK** to proceed with the changes.

Use the online Help if you need additional information for changing your password.

Manage User Profiles and Access

Manage your system users that log on to the HMC. A user profile is a combination of a user ID, server authentication method, permissions, and a text description. Permissions represent the authority levels assigned to the user profile for the objects the user has permission to access.

Users can be authenticated using local authentication on the HMC, by using Kerberos remote authentication, or by using LDAP authentication. For more information on setting up Kerberos authentication on the HMC, see "Manage KDC" on page 86. For more information about LDAP authentication, see "Manage LDAP" on page 86.

For security reasons, remotely authenticated Kerberos or LDAP users cannot lock the local console.

If you are using local authentication, the user ID and password are used to verify a user's authorization to log on the HMC. The user ID must start with an alphabetic character and consist of 1 to 32 characters. The password has the following rules:

- Must begin with an alphanumeric character.
- Must contain at least 8 characters, however, this limit can be changed by your system administrator.
- The characters must be standard 7-bit ASCII characters.
- Valid characters to use for the password can be: A-Z, a-z, 0-9 and special characters (~ ! @ # $ % ^ & * ( ) _ + - = { } [ ] \ : " ; ').

If you are using Kerberos authentication, specify a Kerberos remote user ID.

If you select LDAP authentication, no additional information is required.

The user profile includes managed resource roles and task roles that are assigned to the user. The *managed resource roles* assign permissions for a managed object or group of objects and the *task roles* define the access level for a user to perform on a managed object or group of objects. You can choose from a list of available default managed resource roles, task roles, or customized roles that are created by using the **Manage Task and Resource Roles** task.

See "HMC tasks, user roles, IDs, and associated commands" on page 8 for a listing of all the HMC tasks and the predefined default user IDs that can perform each task.

The default managed resource roles include:

- All System Resources

The default task roles include:

- hmcservicerep (Service Representative)
- hmcviewer (Viewer)
- hmcoperator (Operator)
- hmcpe (Product Engineer)
- hmcsuperadmin (Super Administrator).

To add or customize a user profile, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage User Profiles and Access**.
3. Complete one of the following steps:
   - From the **User Profiles** window, if you are creating a new user ID, point to **User** on the menu bar and when its menu is displayed, click **Add**. The **Add User** window is displayed.
   - From the **User Profiles** window, if you are creating a user ID with the same attributes as an existing profile, point to **User** on the menu bar and when its menu is displayed, click **Copy**. The **Copy User** window is displayed.

     Note: Some user profiles are predefined, such as a default ID, and those permissions cannot be changed. However, you can copy a default user profile, such as operator, and then modify the resulting new user profile. The newly defined user cannot have greater permissions than the original copied user profile.
   - From the **User Profiles** window, if you are deleting a user ID, point to **User** on the menu bar and when its menu is displayed, click **Remove**. The **Remove User** window is displayed.
   - From the **User Profiles** window, if the user ID exists in the window, select the user ID from the list, and then point to **User** on the menu bar and when its menu is displayed, click **Modify**. The **Modify User** window is displayed.
     - To specify timeout and inactivity values, click **User Properties** from the **Modify User** window.
4. Complete or change the fields in the window, click **OK** when you are done.

Use the online Help if you need additional information for creating, modifying, copying, or removing a user profile and modifying timeout and inactivity values.

Adding, Copying, or Modifying User Profiles

Learn how to add, copy, or modify user profiles.

Users who remotely authenticate through Kerberos or Lightweight Directory Access Protocol (LDAP) must have their profiles set appropriately. You must set the user profile of each remotely authenticated Kerberos or LDAP user to use that type of authentication instead of local authentication. A user that is set to use Kerberos or LDAP remote authentication always uses that type of authentication, even when the user logs into the HMC locally.

Note: The use of Kerberos authentication requires configuration of a key distribution center (KDC) server by using the **KDC Configuration** task. The use of LDAP authentication requires configuration of an LDAP server by using the **LDAP Configuration** task. You do not need to set all users to use Kerberos or LDAP remote authentication. You can set some user profiles so that the users can use local authentication only.

From the Adding, Copying, or Modifying User Profiles window, you can modify the following attributes:

- **User ID**: Enter the user ID for the user profile you are creating or managing. The user name must start with an alphabetic character and consist of 1 to 32 characters.
- **Description**: Enter a meaningful description for your own records.
- **Password**: Enter the password for the user ID.
- **Confirm password**: Enter the password again for verification.
- **Password expires in (days)**: Specify the number of days a password is valid before it expires. This input field is available when **Enforce strict password rules** check box is selected.
- **Manage resource roles**: Displays the managed resource roles that are currently available. Select one or more managed resource roles to define access permissions for this user ID.
- **Task roles**: Displays the task roles that are currently available. Select one task role for this user ID.

Use the online Help if you need additional information about creating, modifying, copying, or removing a user profile and modifying timeout and inactivity values.

User Properties

Learn how to specify timeout and inactivity values for the selected user.

You can specify the amount of time for the following timeout and inactivity tasks:

Timeout Values

- **Session timeout minutes**: Specifies the number of minutes during a logon session that a user is prompted for identity verification. If a value other than zero is specified, the user is prompted after the specified time is reached to reenter their password. If a password is not reentered within the specified amount of time in the **Verify timeout minutes** field, the session is disconnected.
- **Verify timeout minutes**: Specifies the amount of time that is required for the user to reenter their password when prompted, if a value was specified in the **Session timeout minutes** field. If the password is not reentered within the specified time, the session is disconnected.
- **Idle timeout minutes**: Specifies the number of minutes the user's session can be idle. If the user does not interact with the session in the specified amount of time, the session is locked and the screen saver starts. Clicking anywhere on the screen prompts the user for identity verification.
- **Minimum time in days between password changes**: Specifies the minimum amount of time in days that must elapse between changes for the user's password.

Note: A note of zero in any of these fields indicates that there is no expiration of time and it is the default value. You can specify up to a maximum value of 525600 minutes (equivalent to one year).

Inactivity Values

- **Disable for inactivity in days**: Specifies the amount of time in days a user is temporarily disabled after the maximum number of days of inactivity is reached.
- **Never disable for inactivity**: Option to never disable a user's session due to inactivity.
- **Allow remote access via the web**: Option to enable remote web server access for the user you are managing.

Manage Users and Tasks

Display the logged on users and the tasks they are running.

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage Users and Tasks**.
3. In the Manage Users and Tasks window, the following information displays:
   - User you are logged in as
   - Time you logged in
   - Number of tasks running

- Your access location
- Information about tasks that are running:
  - Task ID
  - Task name
  - Targets (if any)
  - Session ID

4. Choose to log off or disconnect from a session that is currently running by selecting the session from the Users **Logged On** list, then click **Logoff** or **Disconnect**.

   Alternately, you can choose to switch to another task or end a task by selecting the task from the **Running Tasks** list, then click **Switch To** or **Terminate**.
5. When you have completed this task, click **Close**.

Manage Task and Resource Roles

Use this task to define and customize user roles.

Note: Predefined roles (default roles) cannot be modified.

A *user role* is a collection of authorizations. A user role can be created to define the set of tasks allowed for a given class of user (*task roles*) or it can be created to define the set of managed objects that are manageable for a user (*managed resource roles*). Once you have defined or customized the user roles you can use the **Manage User Profiles and Access** task to create new users with their own permissions.

If the automatic resource role update function is enabled on the Hardware Management Console (HMC) either through the command line interface or through the Rest API CLI runner job, the HMC user can automatically receive permission to the logical partition that is created. If the logical partition is deleted, the permission is automatically revoked.

The predefined managed resource roles include:

- All System Resources

The predefined task roles include:

- hmcservicerep (Service Representative)
- hmcviewer (Viewer)
- hmcoperator (Operator)
- hmcpe (Product Engineer)
- hmcsuperadmin (Super Administrator)

To customize managed resource roles or task roles:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage Task and Resource Roles**.
3. From the **Manage Task and Resource Roles** window, select either **Managed Resource Roles** or **Task Roles**.
4. To add a role, click **Edit** from the menu bar, then click **Add** to create a new role.

   or

   To copy, remove, or modify an existing role, select the object you want to customize, click **Edit** from the menu bar, then click **Copy**, **Remove**, or **Modify**.
5. Click **Exit** when you are have completed the task.

Use the online Help to get additional information for customizing managed resource roles and task roles.

Manage Certificates

Use this task to manage the certificates used on your HMC. It provides the capability of getting information on the certificates used on the console. This task allows you to create a new certificate for the console, change the property values of the certificate, and work with existing and archived certificates or signing certificates.

All remote browser access to the HMC must use Secure Sockets Layer (SSL) encryption. With SSL encryption required for all remote access to the HMC, a certificate is required to provide the keys for this encryption. The HMC provides a self-signed certificate that allows this encryption to occur.

To manage your certificates:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage Certificates**.
3. Use the menu bar from the **Manage Certificates** window for the actions you want to take with the certificates:
   - To create a new certificate for the console, click **Create**, then select **New Certificate**. Determine whether your certificate will be self-signed or signed by a Certificate Authority (CA), then click **OK**.
   - To modify the property values of the self-signed certificate, click **Selected**, then select **Modify**. Make the appropriate changes, then click **OK**.

     Note: If you have a certificate signed by a Certificate Authority (CA) that consists of a root certificate, intermediate certificate, and a client or leaf certificate, complete the following steps to upload the certificate to the HMC:

     - Open the CA signed certificate file by using a text-based editor and split the content of the file and save as three separate files. The first file is the client or leaf certificate, the second file is the intermediate certificate, and the third file is the root certificate.
     - Log in to the HMC to import the certificate. First upload the client certificate and click **Yes** for uploading more files. In the new window, upload the intermediate certificate and the root certificate.
     - Click **OK** to restart the console.
   - To work with existing and archived certificates or signing certificates, click **Advanced**. Then you can choose the following options:
     - Delete existing certificates
     - Work with archived certificates
     - Import certificates
     - View issuer certificates
4. Click **Apply** for all changes to take effect.

Use the online Help if you need additional information for managing your certificates.

Manage Certificate Revocation List

Use this task to create, modify, delete, and import the certificate revocation list that is used on your Hardware Management Console (HMC).

All remote browsers that are accessing the HMC must use Secure Sockets Layer (SSL) encryption. A certificate is required to provide the keys for this encryption. The HMC provides a self-signed certificate that allows this encryption to occur.

To manage your certificate revocation list, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.

2. In the content pane, click **Manage Certificate Revocation List**.
3. Use the menu bar from the **Manage Certificate Revocation List** window for the actions you want to take with the certificates:
   - To create a new certificate revocation list for the console, click **Import**, then select **New CRL**. Determine whether your certification revocation list is imported from removable media on the console or from the file system on the system that is running the web browser.

     Note: If the list is from removable media, then the certificate revocation list file must be in the top directory on the media.
   - To modify a certificate revocation list on the console, select the certification revocation list from the table, and make appropriate changes, then click **Apply**.
   - To delete a certificate revocation list from the console, click **Selected**, then select **Delete CRL**. Select the certification revocation list, then click **OK**.
   - To work with existing and archived certificates or signing certificates, click **Advanced**.

Use the online Help if you need additional information for managing your certificate revocation list.

Manage LDAP

Configure your HMC so that it uses LDAP (Lightweight Directory Access Protocol) authentication.

Before you begin

Note: Before you configure the HMC so that it uses LDAP authentication, you must ensure that a working network connection exists between the HMC and the LDAP servers.

About this task

To configure your HMC so that it uses LDAP authentication, complete the following steps:

Procedure

1. In the navigation area, click **Users and Security**, and then select **Systems and Console Security**.
2. In the content pane, click **Manage LDAP**. The **LDAP Server Definition** window opens.
3. Select **Enable LDAP**.
4. Define an LDAP server to use for authentication (for example, Microsoft Active Directory, Tivoli®, and Open LDAP).
5. Define the LDAP attribute that is used to identify the authenticated user. The default is **uid**, but you can use your own attributes. For Microsoft Active Directory, use **sAMAccountName** as the attribute.
6. Define the distinguished name tree, also known as the search base, for the LDAP server.
7. Click **OK**.

What to do next

If you want to use LDAP authentication, you must configure each remote user's profile so that it uses LDAP remote authentication instead of local authentication.

Manage KDC

View the key distribution center (KDC) servers that are used by this Hardware Management Console (HMC) for Kerberos remote authentication.

From this task, you can complete the following tasks:

- View existing KDC servers.
- Modify existing KDC server parameters that include realm, ticket lifetime, and clock skew.
- Add and configure a KDC server on the HMC.

- Remove a KDC server.
- Import a service key.
- Remove a service key.

Kerberos is a network authentication protocol that is designed to provide strong authentication for client/server applications by using secret-key cryptography.

Under Kerberos, a client (generally either a user or a service) sends a request for a ticket to the KDC. The KDC creates a ticket-granting ticket (TGT) for the client, encrypts it using the client's password as the key, and sends the encrypted TGT back to the client. The client then attempts to decrypt the TGT, by using its password. If the client successfully decrypts the TGT (for example, if the client gave the correct password), it keeps the decrypted TGT, which indicates proof of the client's identity.

The tickets have a time availability period. Kerberos requires the clocks of the involved hosts to be synchronized. If the HMC clock is not synchronized with the clock of KDC server, authentication fails.

A Kerberos realm is an administrative domain, site, or logical network that uses Kerberos remote authentication. Each realm uses a primary Kerberos database that is stored on a KDC server and that contains information about the users and services for that realm. A realm might also have one or more secondary KDC servers, that store read-only copies of the primary Kerberos database for that realm.

To prevent KDC spoofing, the HMC can be configured to use a service key to authenticate to the KDC. Service key files are also known as keytabs. Kerberos verifies that the TGT requested was issued by the same KDC that issued the service key file for the HMC. Before you can import a service key file into an HMC, you must generate a service key for the host principal of the HMC client.

Note: For MIT Kerberos V5 *nix distributions, create a service key file by running the `kadmin` utility on a KDC and by using the `ktadd` command. Other Kerberos implementations might require a different process to create a service key.

You can import a service key file from one of these sources:

- Removable media that is mounted to the HMC, such as optical discs or USB Mass Storage devices. You must use this option locally at the HMC (not remotely), and you must mount the removable media to the HMC before you use this option.
- A remote site that uses secure FTP. You can import a service-key file from any remote site with SSH installed and running.

To use Kerberos remote authentication for this HMC, complete the following tasks:

- You must enable the Network Time Protocol (NTP) service on the HMC and set the HMC and the KDC servers to synchronize time with the same NTP server. You can enable the NTP service on the HMC by accessing the "Change Date and Time" on page 70 task from **Console Management**, and then selecting **Console Settings**.
- You must set the user profile of each remote user to use Kerberos remote authentication instead of local authentication. A user that is set to use Kerberos remote authentication always uses Kerberos remote authentication, even when the user logs on to the HMC locally.

  Note: You do not need to set all users to use Kerberos remote authentication. You can set some user profiles so that the users can use local authentication only.
- Use of a service key file is optional. Before you use a service key file, you must import it into the HMC. If a service key is installed on the HMC, realm names must be equivalent to the network domain name. The following example shows how to create the service key file on a Kerberos server by using the **kadmin.local** command, assuming the HMC hostname is `hmc1`, the DNS domain is `example.com`, and the Kerberos realm name is `EXAMPLE.COM`:

  – `# kadmin_local kadmin.local: ktadd -k /etc/krb5.keytab host/hmc1.example.com@EXAMPLE.COM`

  Using the Kerberos ktutil on the Kerberos server, verify the service key file contents. The output looks like the following example:

  – `# ktutil`

```
ktutil: rkt /etc/krb5.keytab
ktutil: l
slot KVNO Principal
---- ----
---------------------------------------------------------------------
1 9 host/hmc1.example.com@EXAMPLE.COM
2 9 host/hmc1.example.com@EXAMPLE.COM
```

- The HMC Kerberos configuration can be modified for SSH (Secure Shell) login without a password by using GSSAPI. For remote login without a password through Kerberos to an HMC, configure the HMC to use a service key. After the configuration is completed, use `kinit -f principal` to obtain forwardable credentials on a remote Kerberos client machine. You can then enter the following command to log in to the HMC without having to enter a password: `$ ssh -o PreferredAuthentications=gssapi-with-mic user@host`.

To manage the KDC, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Manage KDC** window, select the appropriate task from the available options under the **Actions** menu.
4. When you complete the task, click **OK**.

Use the online Help if you need additional information for Managing KDC.

View KDC Server

Display existing key distribution center (KDC) servers on the Hardware Management Console (HMC).

To view existing KDC Servers on your HMC, click **Users and Security**, and then select **Users and Roles**. In the content pane, click **Configure KDC**. If no servers exist and NTP has not yet been enabled, a warning panel message displays. Enable the NTP service on the HMC and configure a new KDC server as desired.

Modify KDC Server

Learn how to modify the key distribution center (KDC) on your Hardware Management Console (HMC).

To modify existing key distribution center (KDC) server parameters, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. Select a KDC Server.
4. Select a value to modify:
   - **Realm**. A realm is an authentication administrative domain. Normally, realms always appear in upper case letters. It is good practice to create a realm name that is the same as your DNS domain (in upper case letters). A user belongs to a realm if and only if the user shares a key with the authentication server of that realm. Realm names must be equivalent to the network domain name if a service key file is installed on the HMC.
   - **Ticket Lifetime**. Ticket lifetime sets the lifetime for credentials. The format is an integer number followed by one of **s** seconds, **m** minutes, **h** hours, or **d** days. Enter a Kerberos lifetime string such as *2d4h10m*.
   - **Clock skew**. Clock skew sets the maximum allowable amount of clock skew between the HMC and the KDC server before Kerberos considers messages invalid. The format is an integer number that represents number of seconds.
5. Click **OK**.

Add KDC server

Add a Key Distribution Center (KDC) server to this Hardware Management Console (HMC).

To add a new KDC server, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Actions** drop down list, select **Add KDC Server**.
4. Enter the host name or IP address of the KDC server.
5. Enter the KDC server realm.
6. Click **OK**.

Remove KDC server

Kerberos authentication on the Hardware Management Console (HMC) remains enabled until all key distribution center (KDC) servers are removed.

To remove a KDC server:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. Select the KDC server from the list.
4. From the **Actions** drop down list, select **Remove KDC Server**.
5. Click **OK**.

Import Service Key

Before you can import a service key file into an Hardware Management Console (HMC), a service key file must first be created on the Kerberos server for the HMC host. The service key file contains the host principal of the HMC client, for example, `host/example.com@EXAMPLE.COM`. In addition to KDC Authentication, the host service key file is used to enable password-less SSH (Secure Shell) login using GSSAPI.

Note: For MIT Kerberos V5 *nix distributions, create a service key file by running the `kadmin` utility on a KDC and using the `ktadd` command. Other Kerberos implementations may require a different process to create a service key.

To import a service key, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Actions** drop down list, select **Import Service Key**.
4. Select from one of the following:
   - **Local** - The service key must be located on removable media currently mounted on the HMC. You must use this option locally at the HMC (not remotely), and you must mount the removable media to the HMC before using this option. Specify the full path of the service key file on the media.
   - **Remote** - The service key must be located on a remote site available to the HMC via secure FTP. You can import a service key file from any remote site that has SSH (Secure Shell) installed and running. Specify the hostname of the site, a user ID and password for the site, and the full path of the service key file on the remote site.
5. Click **OK**.

Implementation of the service key file will not take effect until the HMC is rebooted.

Remove Service Key

Learn how to remove a service key from your Hardware Management Console (HMC).

To remove the service key from the HMC, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Manage KDC**.
3. From the **Actions** drop down list, select **Remove Service Key**.
4. Click **OK**.

You must reboot the HMC after removing the service key. Failure to reboot may cause login errors.

Manage MFA

Learn how to enable Multi-Factor Authentication (MFA) on the Hardware Management Console (HMC).

Notes:

1. Multi-Factor Authentication is disabled on the HMC by default.
2. For HMC GUI login, when MFA is enabled and the user is configured on the PowerSC MFA server, enter the Cache Token Credential (CTC) code in the password field.
3. For Secure Shell (SSH) login:

   When MFA is enabled, all users that login through SSH are prompted for a CTC code. If the user is configured on the PowerSC MFA server, then you can enter the CTC code at the prompt. If the user is not configured on the PowerSC MFA server, press `Enter` when prompted for CTC code, and then enter the password of the user at the prompt.

To enable Multi-Factor Authentication, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Systems and Console Security**.
2. In the content pane, click **Manage MFA**.
3. From the **Manage MFA** window, select the **Enable multi factor authentication** check box.
4. Enter the following information:
   - **Host name or IP address of the authentication server**
   - **Port of the authentication server**
5. Click **OK**.

Use the online Help if you need additional information about this task.

Enable Remote Command Execution

This task is used to enable remote command execution using the ssh facility.

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Command Execution**.
3. From the **Enable Remote Command Execution** window, select **Enable remote command execution using the ssh facility**.
4. Click **OK**.

Enable Remote Operation

This task is used to allow the HMC to be accessed at a remote workstation through a web browser.

To enable the HMC remote access:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Operation**.
3. Select **Enabled** from the Remote Operation drop-down list, then click **OK**. The HMC can be accessed from a remote workstation using a Web browser.

Use the online Help to get additional information for allowing remote access to the HMC.

Enable Remote Virtual Terminal

A Remote Virtual Terminal connection is a terminal connection to a logical partition from another remote HMC. Use this task to enable Remote Virtual Terminal access for remote clients.

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Virtual Terminal**.
3. From the **Enable Remote Virtual Terminal** window, you can enable this task by selecting Enable remote virtual terminal connections.
4. Click **OK** to activate your changes.

Use the online Help to get additional information for enabling a remote terminal connection.

Serviceability tasks

The tasks that are available on the HMC for the **Serviceability** tasks are described.

Note: Depending on the task roles assigned to your user ID you may not have access to all the tasks. See "HMC tasks, user roles, IDs, and associated commands" on page 8 for a listing of the tasks and the user roles allowed to access them.

Tasks Log

View all the tasks that are currently running or completed on the Hardware Management Console (HMC).

To view the tasks log, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Tasks Log**.
2. You can view the following tabs in the tasks log:
   - **Task name**: Displays the name of task.
   - **Status**: Displays the current state of the task (running or completed).
   - **Resource**: Displays the name of the resource.
   - **Resource type**: Displays the type of resource.
   - **Initiator**: Displays the name of the user that initiated the task.
   - **Start time**: Displays the time that the task was initiated.
   - **Duration**: Displays the amount of time that the task took to complete.

Use the online Help for additional information about viewing the tasks log.

Console Events Logs

View a record of system events occurring on the Hardware Management Console (HMC). System events are individual activities that indicate when processes occur, begin and end, succeed or fail.

To view console events legs, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Console Events Logs**.

2. Use the menu bar to change to a different time range, or to change how the events display in the summary. You can also use the table icons or the **Select Action** menu on the table toolbar to display different variations of the table.
3. When you are done viewing the events, select **View** on the menu bar, then click **Exit**.

Use the online Help for additional information about viewing HMC events.

Serviceable Events Manager

This task allows you to select the criteria for the set of serviceable events you want to view. When you finish selecting the criteria, you can view the serviceable events that match your specified criteria.

To set the criteria for the serviceable events you want to view, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Serviceable Events Manager**.
2. From the **Serviceable Events Manager** window, provide event criteria, error criteria, and FRU criteria.
3. Click **OK** when you have specified the criteria you want for the serviceable events you want to view.

Use the online Help if you need additional information managing events.

Events Manager for Call Home

Learn how to monitor and approve any data that is being transmitted from an HMC to IBM.

1. In the navigation area, click **Serviceability**, and then select **Events Manager for Call Home**.
2. From the **Events Manager for Call Home** window, select **Manage Consoles** to manage the list of registered management consoles. You can use the **Event Criteria** to specify the approval state, status, and originating HMC to filter the list of events that are available for all registered management consoles. You can use the criteria to filter the view and select events to view details, view files, and complete call home operations.
3. Click **OK** to exit Events Manager for Call Home and to save the filter values.

Use the online Help if you need additional information about this task.

Create Serviceable Event

This task reports problems that occurred on your Hardware Management Console (HMC) to the service provider (for example, the mouse does not work) or lets you test problem reporting.

Submitting a problem is dependent upon whether you have customized this Hardware Management Console to use the Remote Support Facility (RSF) and if it is authorized to automatically call for service. If so, the problem information and service request is sent to the service provider automatically with a modem transmission.

To report a problem on your Hardware Management Console, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Create Serviceable Event**.
3. From the **Create Serviceable Event** window, select a problem type from the list displayed.
4. Enter a brief description of your problem in the **Problem Description** input field and then click **Request Service**.

To test problem reporting from the **Report a Problem** window:

1. Select **Test automatic problem reporting** and enter *This is just a test* in the **Problem Description** input field.

2. Click **Request Service**. The problems are reported to the service provider for the Hardware Management Console. Reporting a problem sends to the service provider the information you provide on **Report a Problem** window, and machine information that identifies the console.

Use the online Help if you need additional information for reporting a problem or testing if problem reporting works.

Manage Dumps

Learn how to manage the procedures for dumps of selected systems on your Hardware Management Console (HMC).

To manage a dump, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Dumps**.
3. From the **Manage Dumps** window, select a dump and perform one of the following dump-related tasks:

   From **Selected** on the menu bar:

   - Copy the dump to media.
   - Copy the dump to a remote system.
   - Use the call home feature to transmit the dump to your service provider.
   - Delete a dump.

   From **Actions** on the menu bar:

   - Initiate a dump of the hardware and server firmware for the managed system.
   - Initiate a dump of the service processor.
   - Initiate a dump of the Bulk Power Control service processor.
   - Modify the dump capability parameters for a dump type.

   From **Status** on the menu bar, you can view the offload progress of the dump.
4. Click **OK** when you complete this task.

Use the online Help to get additional information for managing dumps.

Transmit Service Information

Transmit service information to your service provider immediately or schedule when to transmit service information for use for problem determination.

To schedule or transmit service information, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Transmit Service Information**.
3. In the content pane, click the **Schedule and Send Data** tab to schedule the service information.

   Note: You can also click the following tabs to select the data that you want to send and to configure FTP connections:

   - **Schedule and Send Data**: Transmit information to your service provider immediately or schedule the transmission.
   - **Configure FTP Connection**: Provide configuration data to allow the use of FTP to offload service information.
   - **Send Problem Reports**: Select the data that you want and the destination for the data.

4. Select the types of service information that you want to enable regular transmissions or to send immediately.
   - **Operational Test (Heartbeat) Information -- always enabled**: Send the Problem Event log file.
   - **Hardware Service Information (VPD)**: Send the Vital Product Data (VPD) for all managed systems that are attached to this HMC.
   - **Software Service Information**: Send the VPD for all software that is running on the partitions.
   - **Performance Management Information**: Gather and send the performance management information.
   - **Update Access Key Information**: Verifies and updates Access Key information.
5. Select the interval (in days) and the time to schedule repeating transmissions. To transmit the information immediately, click **Send Now**.
6. Click **OK**.

Use the online Help for additional information about scheduling service information.

Format Media

This task formats a diskette or USB 2.0 Flash Drive Memory Key.

You can format a diskette by supplying a user-specified label.

To format a diskette or USB 2.0 Flash Drive Memory Key, complete the following steps:

1. In the navigation area, click **Console Management**, and then select **Console Managment**.
2. In the content pane, click **Format Media**.
3. From the **Format Media** window, select the type of media you want to format, then click **OK**.
4. Make sure that your media is correctly inserted, then click **Format**. The **Format Media** progress window is displayed. When the media is formatted, the **Format Media Completed** window is displayed.
5. Click **OK** and then click **Close** to end the task.

Use the online Help if you need additional information about this task.

Electronic Service Agent Setup Wizard

Learn how to open the Electronic Service Agent Setup wizard using the Hardware Management Console (HMC) interface.

About this task

To open the Electronic Service Agent Setup wizard, complete the following steps:

Procedure

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the contents pane, select **Electronic Service Agent Setup Wizard**. The Electronic Service Agent wizard opens. Follow the instructions in the wizard to configure call-home tasks.

Authorize User

Request authorization for Electronic Service Agent. Electronic Service Agent associates your system with a user ID and allows access to system information through the Electronic Service Agent facility. This registration is also used by your operating system to automate service processes for your AIX or IBM i operation system.

To register a user ID, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Authorize User**.
3. Provide a user ID that is registered with the Electronic Service Agent. If you need a user ID, you can register at the IBM Registration website.
4. Click **OK**.

Use the online Help if you need additional information for registering a customer user ID with the eService website.

Enable Electronic Service Agent

This task allows you enable or disable the call-home state for managed systems.

Note: If Customizable Data Replication is **Enabled** on this HMC (using the **Manage Data Replication** task), the data specified in this task may change depending on automatic replication from other HMCs configured on your network. For more information on data replication, see "Manage Data Replication" on page 75.

By enabling the call-home state for a managed system this causes the console to automatically contact a service center when a serviceable event occurs. When a managed system is disabled, your service representative is not informed of serviceable events.

To manage call-home for the system(s):

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Enable Electronic Service Agent**.
3. From the **Enable Electronic Service Agent** window, select a system or systems you want to enable or disable the call-home state.
4. Click **OK** when you have completed the task.

Use the online Help if you need additional information for enabling the Electronic Service Agent.

Manage Outbound Connectivity

Customize the means for outbound connectivity for the Hardware Management Console (HMC) to use to connect to remote service.

Note: If Customizable Data Replication is **Enabled** on this HMC (using the **Manage Data Replication** task), the data specified in this task may change depending on automatic replication from other HMCs configured on your network. For more information on data replication, see "Manage Data Replication" on page 75.

You can configure this HMC to attempt connections through the local modem, Internet, Internet Virtual Private Network (VPN), or through a remote pass-through system. Remote service is a two-way communication between the HMC and the IBM Service Support System for the purpose of conducting automated service operations. The connection can only be initiated by the HMC. IBM Service Support System cannot and never attempts to initiate a connection to the HMC.

You can configure this HMC to attempt connections through the local modem, Internet, Internet Virtual Private Network (VPN), or through a remote pass-through system. Remote service is a two-way communication between the HMC and your support system for the purpose of conducting automated service operations. The connection can only be initiated by the HMC.

To customize your connectivity information, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Outbound Connectivity**.
3. From the **Manage Outbound Connectivity** window select **Enable local server as call-home server** (a check mark appears) before proceeding with the task.

Note: You must first **Accept** the terms described about the information you provided in this task.

This allows the local HMC to connect to your service provider's remote support facility for call-home requests.

4. The dial information window displays the following tabs for providing input:

   - Local Modem
   - Internet
   - Internet VPN
   - Pass-Through Systems

5. If you want to allow connectivity over a modem, use the **Local Modem** tab, then select **Allow local modem dialing for service** .

   a. If your location requires a prefix to be dialed in order to reach an outside line, click **Modern Configuration** and enter the **Dial prefix** in the **Customize Modem Settings** window required by your location. Click **OK** to accept the setting.

   b. Click **Add** from the **Local Modem** tab page to add a telephone number. When local modem dialing is allowed, there must be at least one telephone number configured.

6. If you want to allow connectivity over the Internet, use the **Internet** tab, then select **Allow an existing internet connection for service**.
7. If you want to configure the use of a VPN over an existing Internet connection to connect from the local HMC to your service provider's remote support facility, use the **Internet VPN** tab.
8. If you want to allow the HMC to use the pass-through systems as configured by the TCP/IP address or host name, use the **Pass-Through Systems** tab.
9. When you complete all the necessary fields, click **OK** to save your changes.

Use the online Help if you need additional information for customizing outbound connectivity information.

Manage Inbound Connectivity

Learn how to allow your service provider to temporarily access your local console, such as the Hardware Management Console (HMC), or the partitions of a managed system.

To manage inbound connectivity, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Inbound Connectivity**.
3. From the **Manage Inbound Connectivity** settings window, you can perform the following tasks:

   - Use the **Remote Service** tab to provide the information necessary to start an attended remote service session.
   - Use the **Call Answer** tab to provide the information necessary to accept incoming calls from your service provider to start an unattended remote service session.

4. Click **OK** to proceed with your selections.

Use the online Help if you need additional information about this task.

Manage Customer Information

This task enables you to customize the customer information for the Hardware Management Console (HMC).

Note: If Customizable Data Replication is *Enabled* on this HMC (using the **Manage Data Replication** task), the data specified in this task may change depending on automatic replication from other HMCs configured on your network. For more information on data replication, see "Manage Data Replication" on page 75.

The **Manage Customer Information** window displays the following tabs for providing input:

- Administrator
- System
- Account

To customize your customer information, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Customer Information**.
3. From the **Manage Customer Information** window, provide the appropriate information on the **Administrator** page.

   Note: Information is required for fields with an asterisk (*).
4. Select the **System** and **Account** tabs from the **Manage Customer Information** window to provide additional information.
5. Click **OK** when you have completed the task.

Use the online Help to get additional information about customizing your account information.

Manage Event Notification

This task adds email addresses that notify you when problem events occur on your system and configures how you want to receive notification of system events from the Electronic Service Agent.

To set up notification, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Event Notification**.
3. From the **Manage Event Notification** window, you can complete the following tasks:
   - Use the **Email** tab to add the email addresses that are notified when problem events occur on your system and when scheduled operations are planned for your system.
   - Use the **SNMP Trap Configuration** tab to specify locations for sending Simple Network Management Protocol (SNMP) trap messages for Hardware Management Console application programming interface events.
4. Click **OK** when you complete this task.

Use the online Help if you need additional information about this task.

Manage Connection Monitoring

Learn how to configure the timers that the connection monitoring uses to detect outages and enables or disables connection monitoring for selected machines.

You can view and, if authorized, change connection monitoring settings by machine. Connection monitoring generates serviceable events when communication problems are detected between the HMC and managed systems. If you disable connection monitoring, no serviceable events are generated for networking problems between the selected machine and this HMC.

To monitor the connections, complete the following steps:

1. In the navigation area, click **Serviceability**, and then select **Service Management**.
2. In the content pane, click **Manage Connection Monitoring**.

3. From the **Manage Connection Monitoring** window, adjust the timer settings, if required, and enable or disable the server.
4. Click **OK** when you have completed the task.

Use the online Help if you need additional information about connection monitoring.

Remote operations

Connect to and use the Hardware Management Console (HMC) remotely.

Remote operations use the GUI used by a local HMC operator or the command line interface (CLI) on the HMC. You can perform operations remotely in the following ways:

- Use a remote HMC.
- Use a web browser to connect to a local HMC.
- Use an HMC remote command line.

The remote HMC is an HMC that is on a different subnet from the service processor, therefore the service processor cannot be auto discovered with IP multicast.

To determine whether to use a remote HMC or web browser that is connected to a local HMC, consider the scope of control that you need. A remote HMC defines a specific set of managed objects that are directly controlled by the remote HMC, while a web browser to a local HMC has control over the same set of managed objects as the local HMC. The communications connectivity and communications speed is an extra consideration. LAN connectivity provides acceptable communications for either a remote HMC or web browser control.

Using a remote HMC

A remote HMC gives the most complete set of functions because it is a complete HMC. Only the process of configuring the managed objects is different from a local HMC.

As a complete HMC, a remote HMC has the same setup and maintenance requirements as a local Hardware Management Console. A remote HMC needs LAN TCP/IP connectivity to each managed object (service processor) that is to be managed; therefore, any customer firewall that might exist between the remote HMC and its managed objects must allow the HMC to service processor communications to occur. A remote HMC might also need communication with another HMC for service and support. Table 10 on page 98 shows the ports that a remote HMC uses for communications.

Table 10. Ports used by a Remote HMC for Communications

Port	Use
udp 9900	HMC to HMC discovery
tcp 9920	HMC to HMC commands

A remote HMC needs connectivity to IBM (or another HMC that has connectivity to IBM) for service and support. The connectivity to IBM might be in the form of access to the internet (through a company firewall).

Performance and the availability of the status information and access to the control functions of the service processor depends on the reliability, availability, and responsiveness of the customer network that interconnects the remote HMC with the managed object. A remote HMC monitors the connection to each service processor and attempts to recover any lost connections and can report those connections that cannot be recovered.

Security for a remote HMC is provided by the HMC user-login procedures in the same way as a local HMC. As with a local HMC, all communication between a remote HMC and each service processor is encrypted. Certificates for secure communications are provided, and can be changed by the user if wanted.

TCP/IP access to the remote HMC is controlled through its internally managed firewall and is limited to HMC-related functions.

Using a web browser

If you need occasional monitoring and control of managed objects that are connected to a single local Hardware Management Console (HMC), use a web browser. An example of using the web browser might be an off-hours monitor from home by an operator or system programmer.

Each HMC contains a web server that can be configured to allow remote access for a specified set of users. If a customer firewall exists between the web browser and the local HMC, the ports must be accessible and the firewall setup to allow incoming requests on these ports. Table 11 on page 99 shows the ports that a web browser needs for communicating with an HMC.

Table 11. Ports that are used by a web browser for communications to the HMC

Port	Use
TCP 443	Secure (HTTPS) remote interface communication
TCP 8443	Secure browser access to web server communication
TCP 9960	Browser applet communication

[1] This port is opened in the HMC firewall when remote access is enabled in HMC Version 7.8.0 and later. This port must also be opened in any firewall that is between the remote client and the HMC.

After an HMC is configured to allow web browser access, a web browser gives an enabled user access to all the configured functions of a local HMC, except those functions that require physical access to the HMC, such as those that use the local diskette or DVD media. The user interface that is presented to the remote web browser user is the same as that of the local HMC and is subject to the same constraints as the local HMC.

The web browser can be connected to the local HMC by using a LAN TCP/IP connection and by using only encrypted (HTTPS) protocols. Logon security for a web browser is provided by the HMC user-login procedures. Certificates for secure communications are provided, and can be changed by the user.

Performance and the availability of the status information and access to the control functions of the managed objects depends on the reliability, availability, and responsiveness of the network that interconnects the web browser with the local HMC. Because there is no direct connection between the web browser and the individual managed objects, the web browser does not monitor the connection to each service processor, does not perform any recovery, and does not report any lost connections. These functions are handled by the local HMC

The web browser system does not require connectivity to IBM for service or support. Maintenance of the browser and system level is the responsibility of the customer.

If the URL of the HMC is specified by using the format https://*xxx.xxx.xxx.xxx* (where *xxx.xxx.xxx.xxx* is the IP address) and Microsoft Internet Explorer is used as the browser, a host name mismatch message is displayed. To avoid this message, a Firefox browser is used or a host name is configured for the HMC, by using the **Change Network Settings** task (see "Change Network Settings" on page 69), and this host name is specified in the URL instead of an IP address. For example, you can use the format https://*host name.domain_name* or https://*host name* (for example, by using https://hmc1.ibm.com or https://hmc1).

Preparing to use the web browser

Perform the necessary steps to prepare to use a web browser to access the Hardware Management Console (HMC).

Before you can use a web browser to access an HMC, you must complete the following tasks:

- Configure the HMC to allow remote control for specified users.

- For LAN-based connections, you must know the TCP/IP address of the HMC to be controlled, and correctly set up any firewall access between the HMC and the web browser.
- Have a valid user ID and password that is assigned by the access administrator for HMC web access.

Web browser requirements

Learn about the requirements your web browser must meet to monitor and control the Hardware Management Console (HMC).

HMC web browser support requires HTML 2.0, JavaScript 1.0, Java™ Virtual Machine (JVM), Java Runtime Environment (JRE) Version 7, and cookie support in browsers that connect to the HMC. Contact your support personnel to assist you in determining whether your browser is configured with a Java Virtual Machine. The web browser must use HTTP 1.1. If you are using a proxy server, HTTP 1.1 must be enabled for the proxy connections. Additionally, pop-up windows must be enabled for all HMCs addressed in the browser if the browser is running with pop-up windows disabled. The following browsers have been tested:

Google Chrome

HMC Version 8.1 supports Google Chrome Version 33.

Microsoft Internet Explorer

HMC Version 8.1 supports Internet Explorer 9.0, Internet Explorer 10.0, and Internet Explorer 11.0.

Note: The performance CEC task is not supported in Internet Explorer 9.0.

- If your browser is configured to use an Internet proxy, then local IP addresses are included in the exception list. For more information, see your network administrator. If you still need to use the proxy to get to the Hardware Management Console, enable Use **HTTP 1.1 through proxy connections** under the **Advanced** tab in your Internet Options window.

Mozilla Firefox

HMC Version 8.1 supports Mozilla Firefox Version 17 and Mozilla Firefox Version 24 Extended Support Release (ESR). Ensure that the JavaScript options to raise or lower windows and to move or resize existing windows are enabled. To enable these options, click the **Content** tab in the browser's Options dialog, click **Advanced** next to the Enable JavaScript option, and then select the Raise or lower windows option and the Move or resize existing windows options. Use these options to easily switch between HMC tasks. For more information about the latest Mozilla Firefox ESR levels, see Security Advisories for Firefox ESR.

Note: The following restrictions apply when you are using Mozilla Firefox while the HMC is in NIST SP 800-131a security mode:

- Mozilla Firefox cannot be used for the remote client.
- The local console cannot be used.

Other web browser considerations

Session cookies need to be enabled in order for ASMI to work when connected to HMC remotely. The ASM proxy code saves session information and uses it.

Internet Explorer

1. Click **Tools** > **Internet Options**.
2. Click the **Privacy** tab and select **Advanced**.
3. Determine whether **Always allow session cookies** is checked.
4. If not checked, select **Override automatic cookie handling** and **Always allow session cookies**.
5. For the First-party Cookies and Third-party Cookies, choose block, prompt, or accept. Prompt is preferred, in which case you are prompted every time that a site tries to write cookies. Some sites need to be allowed to write cookies.

Firefox

1. Click **Tools** > **Options**.

2. Click the **Cookies** Tab.
3. Select **Allow sites to set cookies**.
4. If you want to allow only specific sites, select **Exceptions**, and add this HMC to allow access.

Using the HMC remote command line

An alternative to performing tasks on the HMC graphical user interface is using the command line interface (CLI).

You can use the command line interface in the following situations:

- When consistent results are required. If you must administer several managed systems, you can achieve consistent results by using the command line interface. The command sequence can be stored in scripts and run remotely.
- When automated operations are required. After you develop a consistent way to manage the managed systems, you can automate the operations by starting the scripts from batch-processing applications, such as the **cron** daemon, from other systems.

On a local HMC, you can use the command line interface in the console window.

Setting up secure script execution between SSH clients and the HMC

You must ensure that your script executions between Secure Shell (SSH) clients and the Hardware Management Console (HMC) are secure.

HMCs typically are placed inside the server room where managed systems are located, so you might not have physical access to the HMC. In this case, you can remotely access it using either a remote web browser or the remote command line interface.

Note: To enable scripts to run unattended between an SSH client and an HMC, the SSH protocol must already be installed on the client's operating system.

To enable scripts to run unattended between an SSH client and an HMC, complete the following steps:

1. Enable remote command execution. For more information, see "Enable Remote Command Execution" on page 90.
2. On the client's operating system, run the SSH protocol key generator. To run the SSH protocol key generator, complete the following steps:

   a. To store the keys, create a directory that is named `$HOME/.ssh` (either RSA or DSA keys can be used).

   b. To generate public and private keys, run the following command:

      ```
      ssh-keygen -t rsa
      ```

      The following files are created in the $HOME/.ssh directory:

      ```
      private key: id_rsa
      public key: id_rsa.pub
      ```

      The write bits for both group and other are turned off. Ensure that the private key has a permission of 600.“
3. On the client's operating system, use ssh and run the `mkauthkeys` command to update the HMC user's authorized_keys2 file on the HMC by using the following command:

   ```
   ssh hmcuser@hmchostname mkauthkeys -–add <the contents of $HOME/.ssh/
   id_rsa.pub>
   ```

   Note: Double quotes (“) are used in commands to ensure that the remote shell can properly process the command. For example:

   ```
   ssh "mkauthkeys hscuser@somehmchost --add 'ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDa+Zc8+hn1+
   TjEXu640LqnVNB+UsixIE3c649Cgj20gaVWnFKTjcpWVahK/duCLac/zteMtVAfCx7/ae2g5RTPu7FudF2xjs4r
   +NadVXhoIqmA53a
   ```

```
NjE4GILpfe5vOF25xkBdG9wxigGtJyOKeJHzgnElP7RlEeOBijJDKo5gGE12NVfBxboChm6LtKnDxLi9ahhOYtLlFehJr
6pV/lMAEu
Lhd6ax1hWvwrhf/
h5Ym6J8JbLVL3EeKbCsuG9E4iN1z4HrPkT5OQLqtvC1Ajch1ravsaQqYloMTWNFzM4Qo5O3fZbLc6RuJjtJv8C5t
4/SZUGHZxSPnQmkuii1z9hxt hscpe@vhmccloudvm179'"
```

To delete the key from the HMC, you can use the following command:

```
ssh hmcuser@hmchostname mkauthkeys --remove joe@somehost
```

To enable passwords that prompts for all hosts that access the HMC through SSH, use the `scp` command to copy the key file from the HMC: `scp hmcuser@hmchostname:.ssh/authorized_keys2 authorized_keys2`

Edit the `authorized_keys2` file and remove all lines in this file and then, copy it back to the HMC: `scp authorized_keys2 hmcuser@hmchostname:.ssh/authorized_keys2`

Enabling and disabling HMC remote commands

You can enable or disable the remote command line interface access to the Hardware Management Console (HMC).

To enable or disable remote commands, complete the following steps:

1. In the navigation area, click **Users and Security**, and then select **Users and Roles**.
2. In the content pane, click **Enable Remote Command Execution**.
3. From the **Enable Remote Command Execution** window, select from the following options:
   - To enable remote commands, select **Enable remote command execution using the ssh facility**.
   - To disable remote commands, make sure **Enable remote command execution using the ssh facility** is not selected.
4. Click **OK**.

Logging in to the HMC from a LAN-connected web browser

Log in to the Hardware Management Console (HMC) remotely from a LAN-connected web browser.

Use the following steps to log in to the HMC from a LAN-connected web browser:

1. Ensure that your web browser has LAN connectivity to the wanted HMC.
2. From your web browser, enter the URL of the wanted HMC in the format `https://hostname.domain_name` (for example: `https://hmc1.ibm.com`) or `https://xxx.xxx.xxx.xxx`.

   If this connection is the first access of the HMC for the current web browser session, you might receive a certificate error. This certificate error is displayed if any of the following conditions occur:

   - The web server that is contained in the HMC is configured to use a self-signed certificate and the browser is not configured to trust the HMC as an issuer of certificates.
   - The HMC is configured to use a certificate that is signed by a certificate authority (CA) and the browser is not configured to trust this CA.

   In either case, if you know that the certificate that is being displayed to the browser is the one used by the HMC, you can continue and all communications to the HMC is encrypted.

   If you do not want to receive notification of a certificate error for the first access of any browser session, you can configure the browser to trust the HMC or the CA. In general, to configure the browser, use one of the following methods:

   - You must indicate that the browser permanently trusts the issuer of the certificate.
   - By viewing the certificate and installing to the database of trusted CAs, the certificate of the CA that issues the certificate is used by the HMC.

   If the certificate is self-signed, the HMC itself is considered the CA that issues the certificate.

3. When prompted, enter the user name and password that is assigned by your administrator.

Managing OpenBMC-based and BMC-based systems by using the HMC

Learn how to manage OpenBMC-based and BMC-based systems by using the Hardware Management Console (HMC).

About this task

Learn about the tasks that you perform from the console and how to navigate the baseboard management controller (BMC) by using the web-based user interface with graphical views of managed systems and simplified navigation.

Note: You cannot manage OpenBMC-based and BMC-based systems while the HMC is running in NIST mode.

Add Managed Systems

Learn how to add a managed Baseboard Management Controller (BMC) system to the Hardware Management Console (HMC).

To add one or more managed BMC systems to the HMC, complete the following steps:

1. From the HMC dashboard, click **Connect Systems**
2. From the **Add Managed Systems** window, you can add a BMC system by completing the following fields:

   - **IP Address/Host name**
   - **Username (BMC system)**
   - **Password**

   Alternatively, you can specify a range of IP addresses and click **OK** to view a list of systems that were discovered. You can select one or more discovered systems to add to the HMC.

   Note: The discovery process can take a long time to complete.
3. Click **OK** to add the managed system to the HMC.

Use the online Help if you need additional information about this task.

Systems Management for Servers

Systems Management displays tasks to manage servers. Use these tasks to set up, configure, view status, troubleshoot, and apply solutions for servers.

These tasks are listed when a managed system is selected. The tasks that are listed in the menu pod change as selections are made in the work area.

Operations

Operations contains the tasks for operating managed systems.

Power Off

Shut down the managed system.

Choose from the following options:

Normal power off

The Normal power off mode shuts down the system's operations in a controlled manner. During the shutdown, programs running active jobs are allowed to perform cleanup (end-of-job processing).

Power On

Use the **Power On** task to start a managed system.

Choose from the following options to power on your managed system:

Normal: Select this option to specify that the HMC uses the current setting for the partition start policy to determine how to power on the managed system. The default setting is set to the following value:

- **Auto-Start Always**: This option specifies that the HMC power on logical partitions automatically after the managed system powers on. If powering on the managed system is the result of a user action, the HMC starts all partitions that are configured for automatic startup. If powering on the managed system is the result of an automatic recovery process, the HMC starts only those logical partitions that were running at the time the system is powered off. This option is always available for selection.

Schedule Operations

Create a schedule for certain operations to be performed on the managed system without operator assistance.

Scheduled operations are helpful for situations where automatic, delayed, or repetitious processing of system operations is necessary. A scheduled operation is started at a specified time, without operator assistance to perform the operation. A schedule can be set for one operation or repeated many times.

For example, you might schedule power on or off operations for a managed system.

The Scheduled Operations task displays the following information for each operation:

- The processor that is the object of the operation.
- The scheduled date
- The scheduled time
- The operation
- The number of remaining repetitions

From the **Scheduled Operations** window, you can perform the following tasks:

- Schedule an operation to run later.
- Define operations to repeat at regular intervals.
- Delete a previously scheduled operation.
- View details for a currently scheduled operation.
- View scheduled operations within a specified time range.
- Sort scheduled operations by date, operation, or managed system.

You can schedule an operation to occur once or you can schedule it to repeat. You must provide the time and date that you want the operation to occur. If you want the operation to repeat, you are asked to select the following options:

- The day or days of the week that you want the operation to occur. (optional)
- The interval, or time between each occurrence. (required)
- The total number of repetitions. (required)

The operations that you can schedule for the managed system include the following operations:

Power Off Managed System
Schedules an operation for a system power off at regular intervals for a managed system.

Power On Managed System
Schedules an operation for a system power-on at regular intervals for a managed system.

To schedule operations on the managed system, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. In the content pane, select one or more managed systems.

3. In the menu pod, select **Actions** > **View All Actions** > **Operations** > **Schedule Operations**.
4. From the **Scheduled Operations** window, click **Options** from the menu bar to display the next level of options:
   - To add a scheduled operation, click **Options** and then click **New**.
   - To delete a scheduled operation, select the operation that you want to delete, point to **Options** and then click **Delete**.
   - To update the list of scheduled operations with the current schedules for the selected objects, point to **Options** and then click **Refresh**.
   - To view a scheduled operation, select the operation that you want to view, point to **View** and then click **Schedule Details**.
   - To change the time of a scheduled operation, select the operation that you want to view, point to **View** and then click **New Time Range**.
   - To sort the scheduled operations, point to **Sort** and then click one of the sort categories that appears.
5. To return to the HMC workplace, point to **Operations** and then click **Exit**.

Launch BMC System Management

The Hardware Management Console (HMC) can connect directly to the Baseboard Management Controller (BMC) for a selected system.

The BMC system management is an interface to the service processor that allows you to manage the operation of the server, such as auto power restart, and to view information about the server, such as the error log and vital product data.

To connect to the BMC, complete the following steps:

Note: To access the BMC user interface, you must be at the console or have access to the BMC by using a supported web browser.

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. In the content pane, select one or more managed systems.
3. In the menu pod, select **Actions** > **View All actions** > **Operations** > **Launch BMC System Management**.
4. Click **Continue**.

Configuring Call Home

Problems on your BMC-based managed system are reported to the Hardware Management Console (HMC) as events. You can set up alerts to be automatically notified of any events.

Note: You must enable SNMP traps in the HMC to receive alerts. To enable SNMP traps, navigate to **Console Settings** > **Change Network Settings** > **LAN Adapters** > **Details** > **Firewall Settings**. Select **SNMP Traps** and **SNMP Agent** from the table and then click **Allow Incoming**.

To set up alerts for call home, complete the following steps:

1. From the **BMC System Management** window, click **Configuration** > **Alerts**.
2. Select any alert from the table and click **Modify**.

   Note: You can set up multiple HMCs to receive traps. Duplicate reporting of events by multiple HMCs is possible as duplicate event verification is not performed.
3. Complete the following fields:
   - **Event Severity**
   - **Destination IP**
4. Click **Save**.
5. Verify the new alert in the table.

Use the online Help if you need additional information about this task.

Rebuild

You can extract the configuration information from the managed system and rebuild the information on the Hardware Management Console (HMC).

This task does not disrupt the operation of the running server.

Rebuilding the managed system updates the information on the HMC about the managed system. Rebuilding the managed system is useful when the state of the managed system is `Incomplete`. The Incomplete state means that the HMC cannot gather complete information from the managed system about logical partitions, profiles, or resources.

Rebuilding the managed system is different from refreshing the **HMC** window. When the managed system is rebuilt, the HMC extracts the information from the managed system. You cannot start other tasks while the HMC rebuilds the managed system. This process can take several minutes.

Updates

Display tasks to view system information, manage updates on your Hardware Management Console (HMC), or check system readiness.

Change Licensed Internal Code

Change the Licensed Internal Code of a managed BMC system by using your Hardware Management Console (HMC).

The system firmware is a combination of the BMC firmware and the PNOR firmware. You must update both the BMC firmware and the PNOR firmware for the system to operate properly. If you update only one type of firmware, but do not update the other type of firmware, system errors might occur.

To change the Licensed Internal Code, complete the following steps:

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view system information.
3. In the menu pod, expand **Actions** and then expand **Updates**.
4. Select **Change Licensed Internal Code** > **BMC Change Licensed Internal Code**.
5. Follow the onscreen instructions in the **BMC Change Licensed Internal Code** guided wizard.

   Note: The BMC system must be in the powered off state before you can proceed with the wizard.
6. When you complete this task, click **Close**.

Use the online Help if you need additional information about this task.

Attention LED

View system attention LED information, light specific LEDs to identify a system component, and test all LEDs on a managed system.

The system provides several LEDs that help identify various components, such as enclosures or field replaceable units (FRUs), in the system. For this reason, they are called **Identify** LEDs. Individual LEDs are on or near the components. The LEDs are located either on the component itself or on the carrier of the component (for example, memory card, fan, memory module, or processor). LEDs are either green or amber. Green LEDs indicate either of the following states:

- Electrical power is present.
- Activity is occurring on a link. (The system might be sending or receiving information).

Amber LEDs indicate a fault or identify condition. If your system or one of the components on your system has an amber LED turned on or flashing, identify the problem and take the appropriate action to restore the system back to normal.

You can activate or deactivate the following types of identify LEDs:

Identify LED for an enclosure

If you want to add an adapter to a specific drawer (enclosure), you need to know the machine type, model, and serial number (MTMS) of the drawer. To determine whether you have the correct MTMS for the drawer that needs the new adapter, you can activate the LED for a drawer and verify that the MTMS corresponds to the drawer that requires the new adapter.

You can deactivate a system attention LED. For example, you might determine that a problem is not a high priority and decide to repair the problem later. However, you want to be alerted if another problem occurs, so you must deactivate the system attention LED so that it can be activated again if another problem occurs.

Connections

You can view the Hardware Management Console (HMC) connection status to service processors, reset those connections, connect another HMC to the selected managed system, or disconnect another HMC.

If you select a managed system in the work area, the following tasks pertain to that managed system.

Service Processor Status

View information about the status of the Hardware Management Console (HMC) connection to the service processors on the managed system.

About this task

To show the service processor connection status to the service processors on the managed system, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server for which you want to view service processor connection status.
3. In the menu pod, select **Actions** > **View All Actions** > **Connections** > **Service Processor Status**.

Reset or Remove Connections

Reset or remove a managed system from the Hardware Management Console (HMC) interface.

About this task

To reset or remove connections, complete the following steps:

Procedure

1. In the navigation area, click **Resources**, and then select **All Systems**.
2. Select the server that you want to reset or remove.
3. In the menu pod, select **Actions** > **View All Actions** > **Connections** > **Reset or Remove Connections**.
4. Select **Reset Connection** or **Remove Connection**.
5. Click **OK**.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Accessibility features for IBM Power servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with ICT Accessibility 508 Standards and 255 Guidelines (https://www.access-board.gov/ict/) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power servers.

The IBM Power servers online product documentation in IBM Documentation is enabled for accessibility. For more information about IBM's commitment to accessibility, see the IBM accessibility website at IBM Accessibility (https://www.ibm.com/able/).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s user name and IP address for purposes of session management. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://

www.ibm.com/privacy/details the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Programming interface information

This Managing the Hardware Management Console publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM Hardware Management Console Version 10 Release 1 Maintenance Level 1010.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Power Systems

Problem analysis, system parts, and locations for the IBM Power systems HMC (7063-CR2)

Note

Before using this information and the product it supports, read the information in “Safety notices” on page v, “Notices” on page 23, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Power® servers that contain the Power10 processor and to all associated models.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.

- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).
- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.

- For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer's DC power source, when directed to disconnect power during servicing.

- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

⚠ **CAUTION:**

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position (compliance ID RACK-001 or 22U (compliance ID RR001) and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.

  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U (compliance ID RACK-001 or 22U (compliance ID RR001) level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

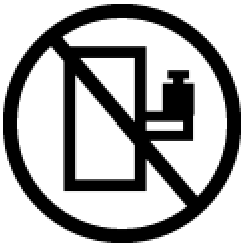

⚠ **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.

- Before extending the rack to the installation position, read the installation instructions.
- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

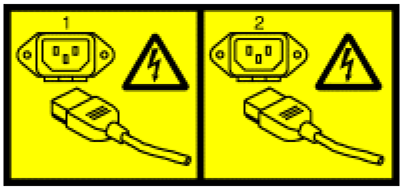

or

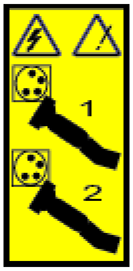

or

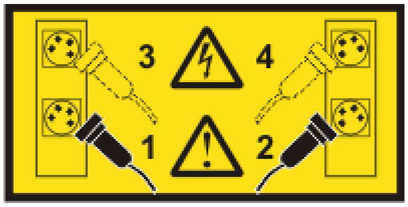

or

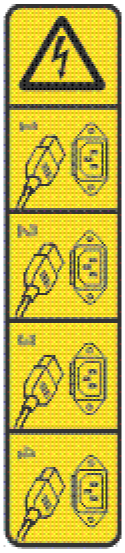

or

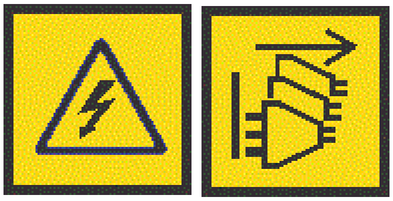

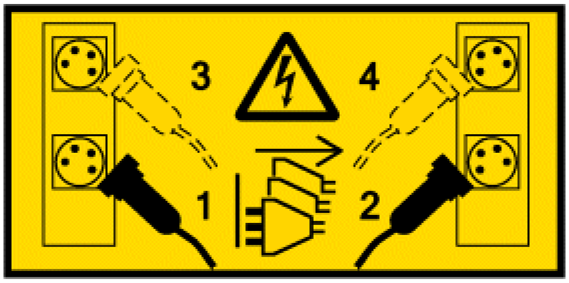

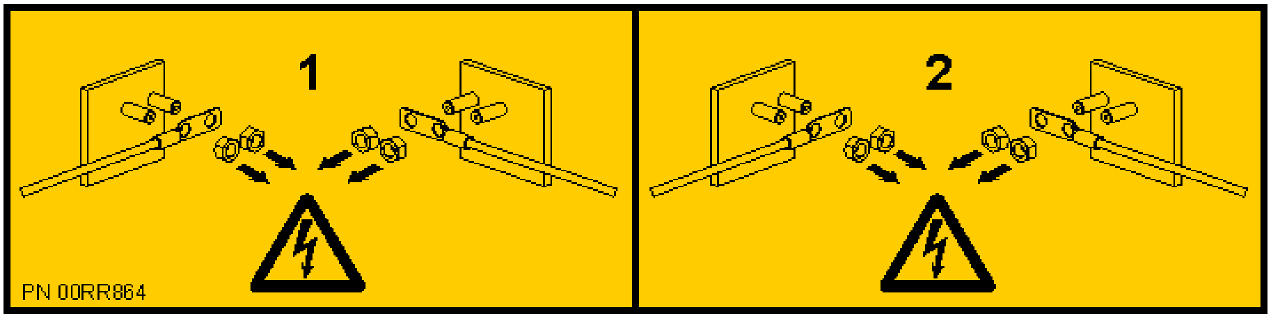

DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

CAUTION: A hot surface nearby. (L007)

(L008)

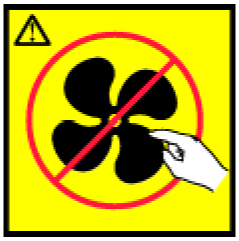

CAUTION: Hazardous moving parts nearby. (L008)

(L018)

or

CAUTION: High levels of acoustical noise are (or could be under certain circumstances) present. Use approved hearing protection and/ or provide mitigation or limit exposure. (L018)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

CAUTION: This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

CAUTION: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

CAUTION: This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

CAUTION: Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

CAUTION: The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approvedapproved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)(C003a)

CAUTION: Regarding IBM providedprovided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.

- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four (4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.
- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.
- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intra-building ports of this equipment are suitable for connection to intra-building or unexposed wiring or cabling only. The intra-building ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The AC-powered system does not require the use of an external surge protection device (SPD).

The DC-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The DC-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Beginning troubleshooting and problem analysis

This information provides a starting point for analyzing problems.

This information is the starting point for diagnosing and repairing systems. From this point, you are guided to the appropriate information to help you diagnose problems, determine the appropriate repair action, and then complete the necessary steps to repair the system.

Notes:

- Update the system firmware to the latest level before you start problem analysis. If you update the system firmware, you have the latest available fixes and improvements for error handling, reporting, and isolation. For instructions about updating the system firmware, see Getting fixes.
- Some service procedures use OpenBMC tool commands. To download and install the OpenBMC tool, see Downloading and installing the OpenBMC tool.

What type of problem are you dealing with?	Problem analysis procedure
You do not know the type of problem.	Go to "Determining the problem analysis procedure to perform" on page 1.
You have an FQPSP*xxxxxxx* event code.	Go to FQPSP*xxxxxxx* Event Codes.
A baseboard management controller (BMC) access problem occurred.	Go to "Resolving a BMC access problem" on page 2.
The system does not power on (the power button or the BMC power on command does not power on the system).	Go to "Resolving a power problem" on page 4.
A system firmware boot failure occurred (the system started but was not able to boot to the Petitboot menu).	Go to "Resolving a system firmware boot failure" on page 5.
A video graphics array (VGA) monitor problem occurred (the system started but no video is displayed on the monitor).	Go to "Resolving a VGA monitor problem" on page 6.
An operating system boot failure occurred (the system booted to the Petitboot menu but the operating system did not start).	Go to "Resolving an operating system boot failure" on page 6.
A processor, memory, power, or cooling hardware failure occurred.	Go to "Resolving a hardware problem" on page 7.
Missing or faulty PCIe adapter or device.	Go to "Resolving a PCIe adapter or device problem" on page 7.

Determining the problem analysis procedure to perform

Learn how to identify the correct problem analysis procedure to perform.

About this task

To determine the correct problem analysis procedure to perform, complete the following steps:

Procedure

1. After you apply power to the system, are the power supply LEDs green?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a power problem" on page 4.

2. Can you access the baseboard management controller (BMC) across the network?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a BMC access problem" on page 2.

3. Can you boot the system to the Petitboot menu?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a system firmware boot failure" on page 5.

4. Is video displayed on the video graphics array (VGA) monitor?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving a VGA monitor problem" on page 6.

5. Can you start the operating system?

If	Then
Yes:	Continue with the next step.
No:	Go to "Resolving an operating system boot failure" on page 6.

6. Go to "Resolving a hardware problem" on page 7. **This ends the procedure.**

Resolving a BMC access problem

Learn how to identify the service action that is needed to resolve a baseboard management controller (BMC) access problem.

Procedure

1. Ensure that the BMC password is not set to the default password. For information about changing the default password, see Logging on to the OpenBMC GUI. Does the problem persist?

If	Then
Yes:	Continue with the next step.
No:	**This ends the procedure.**

2. Are both ends of the network cable seated securely?

If	Then
Yes:	Continue with the next step.
No:	Seat both ends of the cable securely. If the problem persists, continue with the next step.

3. Is the operating system available?

If	Then
Yes:	Continue with the next step.
No:	Continue with step "5" on page 3.

4. Verify that the BMC network settings are correct.
   a) In the navigation area, click **HMC Management**, and then select **Console Settings**.
   b) In the content pane, click **Change BMC/IPMI network settings**.
   c) Verify that the MAC address and the IP address settings are correct.

   Does the BMC access problem persist?

If	Then
Yes:	Continue with the next step.
No:	**This ends the procedure.**

5. Power off the system and disconnect all AC power cords for 30 seconds. Then, reconnect the AC power cords and power on the system. Does the BMC access problem persist?

If	Then
Yes:	Continue with the next step.
No:	**This ends the procedure.**

6. Verify that the BMC network settings are correct.

   Note: To verify the BMC network settings, you must have a cabled serial connection or a monitor and keyboard.

   a) Power on the system by using the power button on the front of the system. Wait 1 - 2 minutes for the system to display the Petitboot menu.
   b) When the Petitboot menu is displayed, press any key to interrupt the boot process. Then, select Exit to Shell.
   c) For a shared BMC Ethernet port, type the following command and press Enter:

```
ipmitool lan print 1
```

   For a dedicated BMC Ethernet port, type the following command and press Enter:

```
ipmitool lan print 2
```

   To determine the location of the shared and dedicated BMC Ethernet ports, see Table 1 on page 3.

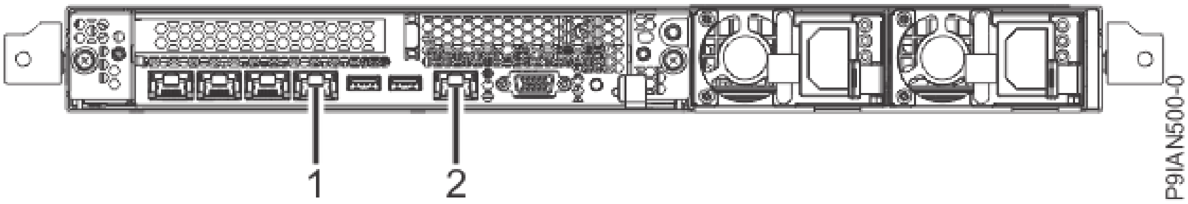

Figure 1. Rear BMC Ethernet ports

Table 1. BMC Ethernet ports

Identifier	Description
1	Shared BMC Ethernet
2	Dedicated BMC Ethernet

   d) Verify that the MAC address and the IP address settings are correct. Then, continue with the next step.

   Note: If the IP address setting is incorrect, go to Configuring the BMC IP address. If the MAC address is 00:00:00:00:00:00, go to “Contacting IBM service and support” on page 12.

7. Complete the following actions:

a. Power on to the Petitboot menu.

b. Update the system firmware. For instructions, see Getting fixes.

Are you able to access the BMC?

If	Then
Yes:	**This ends the procedure.**
No:	Continue with the next step.

8. Replace the system backplane. Go to "7063-CR2 locations" on page 13 to identify the physical location and the removal and replacement procedure. **This ends the procedure.**

Resolving a power problem

Learn how to identify the service action that is needed to resolve a power problem.

Procedure

1. Is the amber LED (bottom LED) of a power supply on solid and is the amber LED on the front of the system turned off?

If	Then
Yes:	Ensure that the power cords for both power supplies are fully seated and that the power distribution units (PDUs) and power outlets are supplying electricity. **This ends the procedure.**
No:	Continue with the next step.

2. Are the power supply LEDs turned off?

If	Then
Yes:	Continue with the next step.
No:	Continue with step "4" on page 4.

3. Perform the following actions, one at a time until the problem is resolved:

   a. Ensure that all of the power cords are fully seated in the power supplies.

   b. Ensure that all of the power cords are fully seated in the power distribution units (PDUs) or wall outlets.

   c. If the power cords are plugged into PDUs, ensure that the PDUs are turned on.

   d. Ensure that all of the power cords are plugged into PDUs or wall outlets that are supplying electricity.

   e. Replace the power cords.

   f. Replace the power supplies. Go to "7063-CR2 locations" on page 13 to identify the physical location and the removal and replacement procedure. **This ends the procedure**.

4. Is the amber LED of a power supply on solid and is the amber system attention LED on the front of the system on solid?

If	Then
Yes:	Continue with the next step.
No:	Go to "Contacting IBM service and support" on page 12. **This ends the procedure.**

5. Perform the following actions, one at a time until the problem is resolved:

   a. Resolve any serviceable alerts that are in the event log. Go to "Resolving a hardware problem" on page 7.

   b. Ensure that the power supply is fully seated in the system.

c. Ensure that the power supply fan is not blocked.

d. Replace the power supply. Go to “7063-CR2 locations” on page 13 to identify the physical location and the removal and replacement procedure. **This ends the procedure**.

Resolving a system firmware boot failure

Learn how to identify the service action that is needed to resolve a failure while booting your system firmware.

Procedure

1. After you press the power button, did the system turn on but fail to display the Petitboot menu?

If	Then
Yes:	Continue with the next step.
No:	Continue with step “6” on page 5.

2. Does the baseboard management controller (BMC) respond to commands?

   Note: To determine whether the BMC responds to commands, run the following OpenBMC tool command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis power status
```

If	Then
Yes:	Continue with the next step.
No:	Continue with step “4” on page 5.

3. Complete the following actions:

   a. Update the system firmware. For instructions, see Getting fixes.

   b. Check the system event logs. For instructions, see “Identifying a service action by using system event logs” on page 10. Then, continue with step “6” on page 5.

4. Disconnect the power cords from the system for 30 seconds. Reconnect the power cords, wait 5 minutes, and then continue with the next step.

5. Does the baseboard management controller (BMC) respond to commands?

   Note: To determine whether the BMC responds to commands, run the following OpenBMC tool command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis power status
```

If	Then
Yes:	Update the system firmware. For instructions, see Getting fixes. **This ends the procedure.**
No:	Replace the system backplane. Go to “7063-CR2 locations” on page 13 to identify the physical location and the removal and replacement procedure. **This ends the procedure.**

6. Power off the system and disconnect all AC power cords for 30 seconds. Then, reconnect the AC power cords and power on the system. Does the system boot successfully?

If	Then
Yes:	**This ends the procedure.**
No:	Go to “Resolving a hardware problem” on page 7. **This ends the procedure.**

Resolving a VGA monitor problem

Learn how to identify the service action that is needed to resolve a video graphics array (VGA) monitor problem.

Procedure

1. Is the system powered on and is the VGA monitor connected to the VGA display port, but no video is displayed?

If	Then
Yes:	Continue with the next step.
No:	Power on the system. **This ends the procedure.**

2. Complete the following steps, one at a time until the problem is resolved:
   a) Ensure that the network image that is specified is a valid boot image.
   b) Ensure that the VGA cable is properly seated to the server port and to the monitor port.
   c) Verify that your monitor and your VGA cable are working properly by testing them on a system that is known to be working properly. If the monitor or the VGA cable does not work properly, replace it.
   d) Verify that the system is powered on by activating a serial over LAN (SOL) session through the baseboard management controller (BMC). If the system is not active, go to "Resolving a system firmware boot failure" on page 5.
   e) Replace the system backplane. Go to "7063-CR2 locations" on page 13 to identify the physical location and the removal and replacement procedure. **This ends the procedure**.

Resolving an operating system boot failure

Learn how to identify the service action that is needed to resolve a failure while booting your operating system.

Procedure

1. Was the system recently installed, serviced, moved, or upgraded?

If	Then
Yes:	Ensure that all cables are properly seated in the connection path to the designated boot device. **This ends the procedure.**
No:	Continue with the next step.

2. Petitboot displays all recognized bootable images to use by default. Is the boot image recognized by Petitboot?

If	Then
Yes:	Continue with the next step.
No:	Select the Petitboot menu option to refresh the boot images. If the problem persists, go to "Resolving a storage device problem" on page 9. **This ends the procedure.**

3. Does an operating system error occur during the boot?

If	Then
Yes:	Recover the operating system with the tools provided for the operating system. If that does not resolve the problem, reinstall the operating system. **This ends the procedure.**
No:	Reinstall the operating system. **This ends the procedure.**

Resolving a hardware problem

Learn how to identify the service action that is needed to resolve a hardware problem.

Procedure

1. If you have not already done so, manually boot the system.
2. Go to "Identifying a service action by using system event logs" on page 10. Then, continue with the next step.
3. Was a service action identified?

If	Then
Yes:	Continue with the next step.
No:	Go to step "5" on page 7.

4. Did the service action fix the problem?

If	Then
Yes:	**This ends the procedure.**
No:	Go to step "5" on page 7.

5. Go to "Resolving a PCIe adapter or device problem" on page 7. Then, continue with the next step.
6. Was a service action identified?

If	Then
Yes:	Continue with the next step.
No:	Go to "Collecting diagnostic data" on page 11. Then, go to "Contacting IBM service and support" on page 12. **This ends the procedure.**

7. Did the service action fix the problem?

If	Then
Yes:	**This ends the procedure.**
No:	Go to "Collecting diagnostic data" on page 11. Then, go to "Contacting IBM service and support" on page 12. **This ends the procedure.**

Resolving a PCIe adapter or device problem

Learn how to access log files, information to identify types of events, and a list of potential problems and service actions.

Procedure

1. To identify the correct service procedure to perform by using operating system log information, complete the following steps:
   a) Log in as the **hscroot** user.
   b) To display the operating system logs, type `less /var/log/messages` and press **Enter**.
2. Scan the operating system logs that occurred around the time that the problem started for the first occurrence of keywords, such as fail, failure, or failed. When you find a keyword that accompanies one or more of the resource names in the following table, a service action is required. Use the following table to determine the service procedure to perform for your type of problem.

Table 2. Resource names, examples, and service procedures for different types of operating system logs.

Resource name	Example of a log requiring a service action	Type of problem	Service procedure
eth1, eth2, eth3, enP*xxxxx*, where *xxxxx* indicates the network port.	`Failed to re-initialize device`	Network	Go to "Resolving a network adapter problem" on page 8.

tg3	`PCI I/O error detected.`
	`Link is Down`
		Network	Go to "Resolving a network adapter problem" on page 8.

sda, sdb, sdc	`FAILED Result`	Storage	Go to "Resolving a storage device problem" on page 9.
EEH	`Detected error on PHB#`*xxx*, where *xxx* is the PHB number.	PCIe bus or adapter	Resolve any device driver errors that are related to I/O and that occurred near the time of this operating system log entry.
	xxx `has failed 6 times in the last hour and has been permanently disabled`, where *xxx* is the PCI bus number.	PCIe bus or adapter	Ensure that the correct device drivers are properly installed for the device. If the problem persists, replace the adapter in the PCIe slot that is specified in the operating system log entry.

Resolving a network adapter problem

Learn about the possible problems and service actions that you can perform to resolve a network adapter problem.

About this task

Table 3. Network adapter problems and service actions

Problem	Service action
System unable to find adapter	1. Verify that the most recent firmware is installed on the system. Otherwise, install the most recent firmware if it is not already installed. 2. Restart the system. 3. Replace the adapter. 4. Replace the system backplane. 5. Replace the central processing unit (CPU).

Table 3. Network adapter problems and service actions (continued)

Problem	Service action
Adapter stops working suddenly	1. If the system was recently installed, moved, serviced, or upgraded, verify that the adapter is seated properly and all associated cables are correctly connected. 2. Inspect the PCIe socket and verify that there is no dirt or debris in the socket. 3. Inspect the card and verify that it is not physically damaged. 4. Verify that all cables are properly seated and are not physically damaged. 5. Replace the adapter. 6. Replace the system backplane. 7. Replace the CPU.
Link indicator light on the adapter is off	1. Verify that the cable functions properly by testing it with a known working connection. 2. Verify that the port or ports on the switch are enabled and functional. 3. Verify that the switch and adapter are compatible. 4. Replace the adapter.
Link light on the adapter is on, but there is no communication from the adapter	1. Verify that the most recent driver is installed, or install the most recent driver if it is not already installed. 2. Verify that the adapter and its link have compatible settings, such as speed and duplex configuration.

Resolving a storage device problem

Learn about the possible problems and service actions that you can perform to resolve a storage device problem.

About this task

Note: To determine the location of the storage device, see Removing and replacing a drive in the 7063-CR2.

Table 4. Storage device problems and service actions

Problem	Service action
System unable to find a storage device that is at the front of the system	1. If the system was recently installed, moved, serviced, or upgraded, verify that the device is seated and installed properly. 2. Verify that the device is compatible with your system. 3. Verify that all internal cables are properly seated and are not physically damaged. 4. Verify that the most recent firmware is installed on the system. Otherwise, install the most recent firmware if it is not already installed. 5. Replace the drive. 6. Replace the cable. 7. Replace the drive holder.
Drive stops working suddenly	1. Verify that all internal cables are properly seated and are not physically damaged. 2. Check the system logs to verify whether the system detected a problem. 3. Replace the drive. 4. Replace the cable.
Other problems	Check the messages and resolve any other problems that were detected. Then, test the drive again. If the drive continues not to function, refer to the documentation for the drive.

Identifying a service action by using system event logs

Use the OpenBMC tool to examine system event logs (SELs) to identify a service action.

Procedure

1. From a system that has the OpenBMC tool installed, type the following command and press Enter:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel print
```

2. Is there an **Active Alerts** section displayed in the output of the command?

If	Then
Yes:	Continue with the next step.
No	No service action is required. **This ends the procedure.** **Note:** Alerts that are displayed in the **Historical Alerts** section do not require service.

3. Is there an entry in the **Active Alerts** section with a value of **Yes** in the **Serviceable** column?

If	Then
Yes:	Continue with the next step.
No	No service action is required. **This ends the procedure.** **Note:** Alerts with a value of **No** that are displayed in the **Serviceable** column do not require service.

4. Starting with the first entry in the **Active Alerts** section with a value of **Yes** in the **Serviceable** column, complete the following steps until all entries are resolved:

   a. Record the log number that is displayed in the **Entry** column.

   b. Record the FQPSP*xxxxxxx* value that is displayed in the **ID** column. Then, go to FQPSP*xxxxxxx* Event Codes and complete the service action that is indicated for the FQPSP*xxxxxxx* event code.

   c. After the service action is complete and the problem is resolved, type the following command and press Enter:

   `openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> sel resolve -n` *x*, where *x* is the log number that you recorded in step "4.a" on page 11.

This ends the procedure.

Verifying a repair

Learn how to verify hardware operation after you make repairs to the system.

Procedure

1. Power on the system.
2. Scan the system event logs (SELs) for serviceable events that occurred after system hardware was replaced. For information about SELs that require a service action, see "Identifying a service action by using system event logs" on page 10.
3. Did any serviceable SEL events occur after hardware was replaced?

If	Then
Yes:	The problem is not resolved. Go to "Identifying a service action by using system event logs" on page 10 and complete the service actions indicated. **This ends the procedure.**
No:	The problem is resolved. **This ends the procedure.**

Collecting diagnostic data

Learn how to collect diagnostic data to send to IBM service and support.

About this task

To collect diagnostic data, complete the following steps:

Procedure

1. Are you able to log on to the Hardware Management Console (HMC)?

If	Then
Yes:	Continue with the next step.
No:	Go to step "3" on page 11.

2. Collect diagnostic data from the 7063-CR2 by using the **`PEDBG`** command on the HMC. To collect diagnostic data, go to HMC Enhanced View: Collecting PEDBG from the HMC (http://www.ibm.com/support/pages/hmc-enhanced-view-collecting-pedbg-hmc) and complete the steps that are indicated. Send the data that you collected during this procedure to IBM service and support. **This ends the procedure.**
3. Can you boot the system to the Petitboot menu or is another system available that has the Linux® operating system?

If	Then
Yes:	Continue with the next step.
No:	Go to "Contacting IBM service and support" on page 12.

4. To collect system event logs, complete the following steps:
   a) Go to the IBM Support Portal (http://www.ibm.com/mysupport/s/).
   b) In the search field, type `Scale-out LC System Event Log Collection Tool`.
   c) Click the **Scale-out LC System Event Log Collection Tool** entry and follow the instructions to install and run the system event log collection tool. Then, continue with the next step.
5. Send the data that you collected during this procedure to IBM service and support. **This ends the procedure.**

Contacting IBM service and support

You can contact IBM service and support by telephone or through the IBM Support Portal.

Before you contact IBM service and support, go to "Beginning troubleshooting and problem analysis" on page 1 and complete all of the service actions indicated. If the service actions do not resolve the problem, or if you are directed to contact support, go to "Collecting diagnostic data" on page 11. Then, use the information below to contact IBM service and support.

Customers in the United States, United States territories, or Canada can place a hardware service request online. To place a hardware service request online, go to the IBM Support Portal (http://www.ibm.com/support/entry/portal/product/power/scale-out_lc).

For up-to-date telephone contact information, go to the Directory of worldwide contacts website (www.ibm.com/planetwide/).

Table 5. Service and support contacts

Type of problem	Call
• Advice • Migrating • "How to" • Operating • Configuring • Ordering • Performance • General information	• 1-800-IBM-CALL (1–800–426–2255) • 1-800-IBM-4YOU (1–800–426–4968)
Software: • Fix information • Operating system problem • IBM application program • Loop, hang, or message Hardware: • IBM system hardware broken • Hardware reference code • IBM input/output (I/O) problem • Upgrade	1-800-IBM-SERV (1–800–426–7378)

Finding parts and locations

Locate physical part locations and identify parts with system diagrams.

Locate the FRU

Use the graphics and tables to locate the field-replaceable unit (FRU) and identify the FRU part number.

7063-CR2 locations

Use this information to find the location of a field-replaceable unit (FRU) in the system unit.

Rack views

The following diagrams show FRU layouts in the system. Use these diagrams with the following tables.

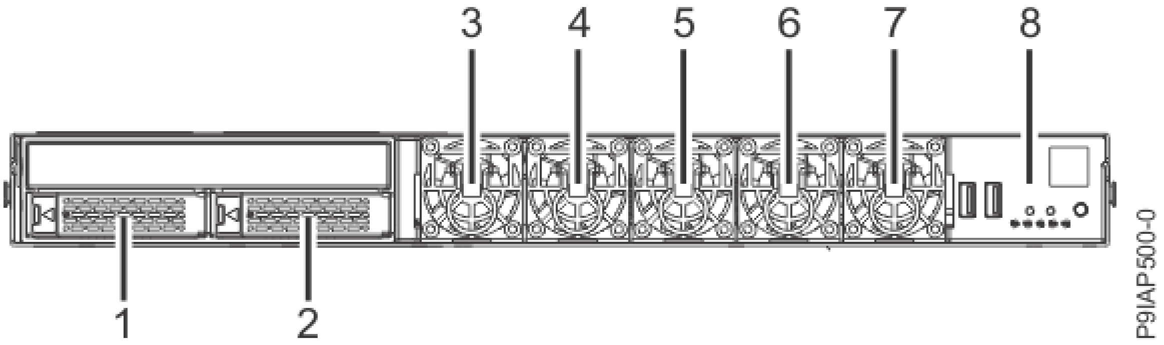

Figure 2. Front view

Table 6. Front view locations

Index number	FRU description	FRU removal and replacement procedures
1	Drive 0	See Removing and replacing a drive in the 7063-CR2.
2	Drive 1	
3	Fan 0	See Removing and replacing fans in the 7063-CR2.
4	Fan 1	
5	Fan 2	
6	Fan 3	
7	Fan 4	
8	Control panel	See Removing and replacing the control panel in the 7063-CR2.

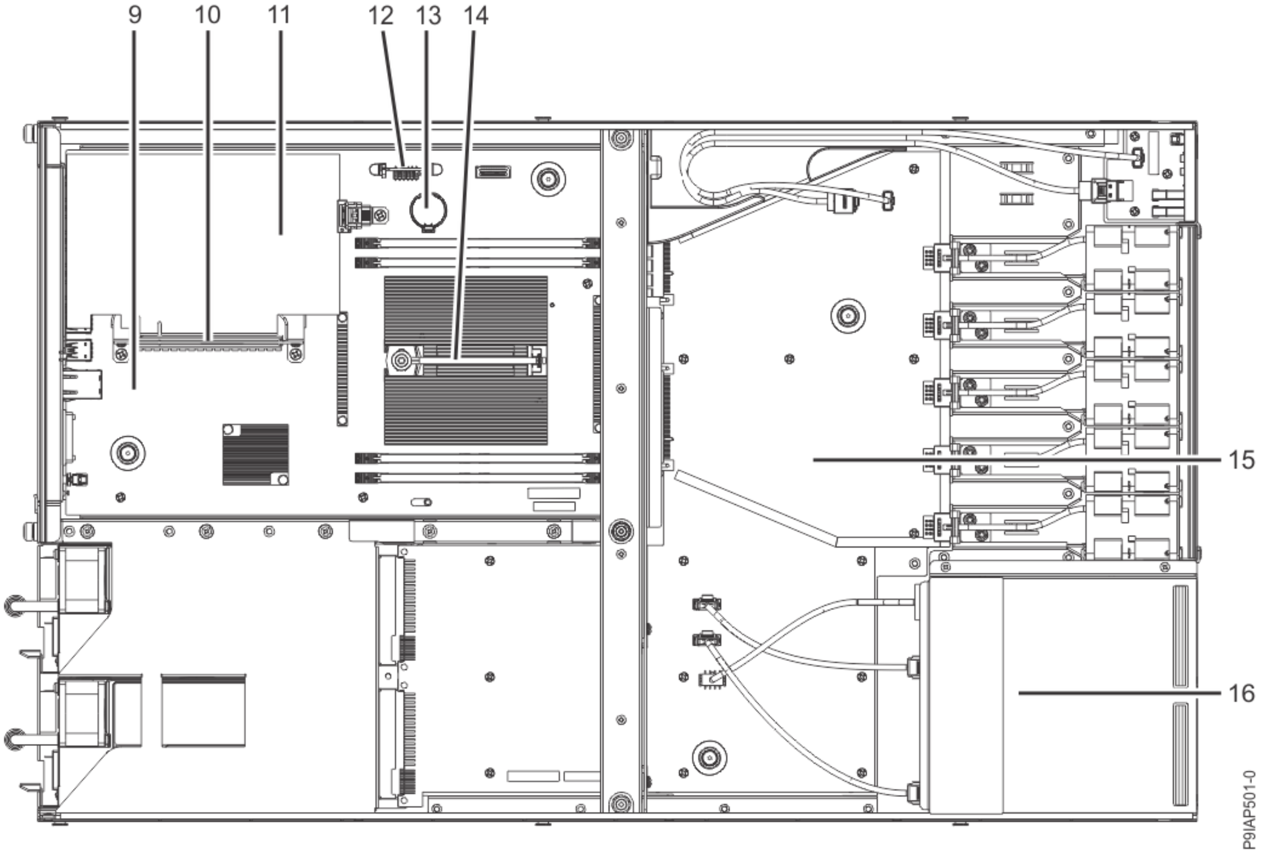

Figure 3. Top view

Table 7. Top view locations		
Index number	**FRU description**	**FRU removal and replacement procedures**
9	System backplane	See Removing and replacing the system backplane in the 7063-CR2.
10	Riser	See Removing and replacing the PCIe riser in the 7063-CR2.
11	PCIe adapter	See Removing and replacing PCIe adapters in the 7063-CR2.
12	Trusted platform module	See Removing and replacing the trusted platform module in the 7063-CR2.
13	Time-of-day battery	See Removing and replacing the time-of-day battery in the 7063-CR2.
14	CPU 0	See Removing and replacing the system processor module in the 7063-CR2.
15	Power distribution board	See Removing and replacing the power distribution board in the 7063-CR2.

Table 7. Top view locations (continued)

Index number	FRU description	FRU removal and replacement procedures
16	Drive holder	See Removing and replacing the drive holder in the 7063-CR2.

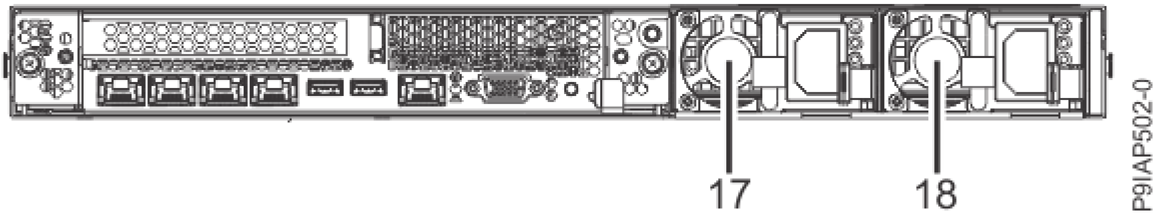

Figure 4. Rear view

Table 8. Rear view locations

Index number	FRU description	FRU removal and replacement procedures
17	Power supply unit 1 (PSU 1)*	See Removing and replacing a power supply in the 7063-CR2.
18	Power supply unit 0 (PSU 0)*	

* The E*x* labels on the chassis do not match the power supply unit number.

Memory locations

The following diagram shows memory DIMMs and their corresponding field-replaceable unit (FRU) layouts in the system. Use this diagram with the following table.

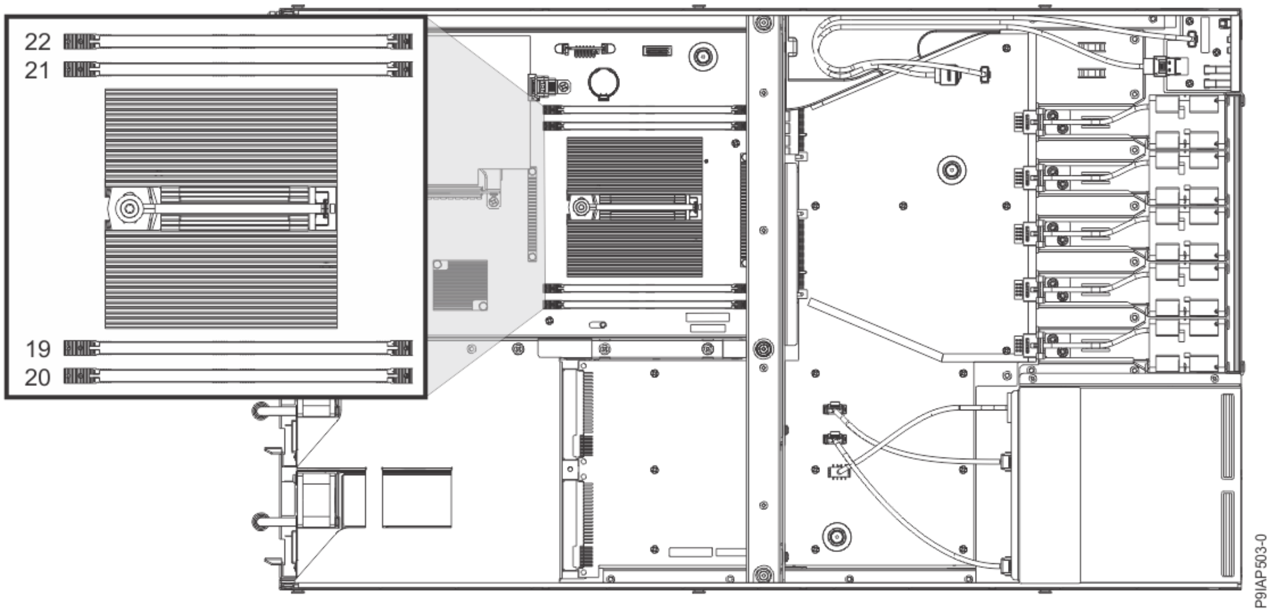

Figure 5. Memory locations

Table 9. Memory locations

Index number	FRU description	FRU removal and replacement procedures
19	DIMM 0	See Removing and replacing memory in the 7063-CR2.
20	DIMM 1	
21	DIMM 2	
22	DIMM 3	

7063-CR2 parts

Use this information to find the field-replaceable unit (FRU) part number.

After you identify the part number of the part that you want to order, go to Advanced Part Exchange Warranty Service. Registration is required. If you are not able to identify the part number, go to Contacting IBM service and support.

Rack final assembly

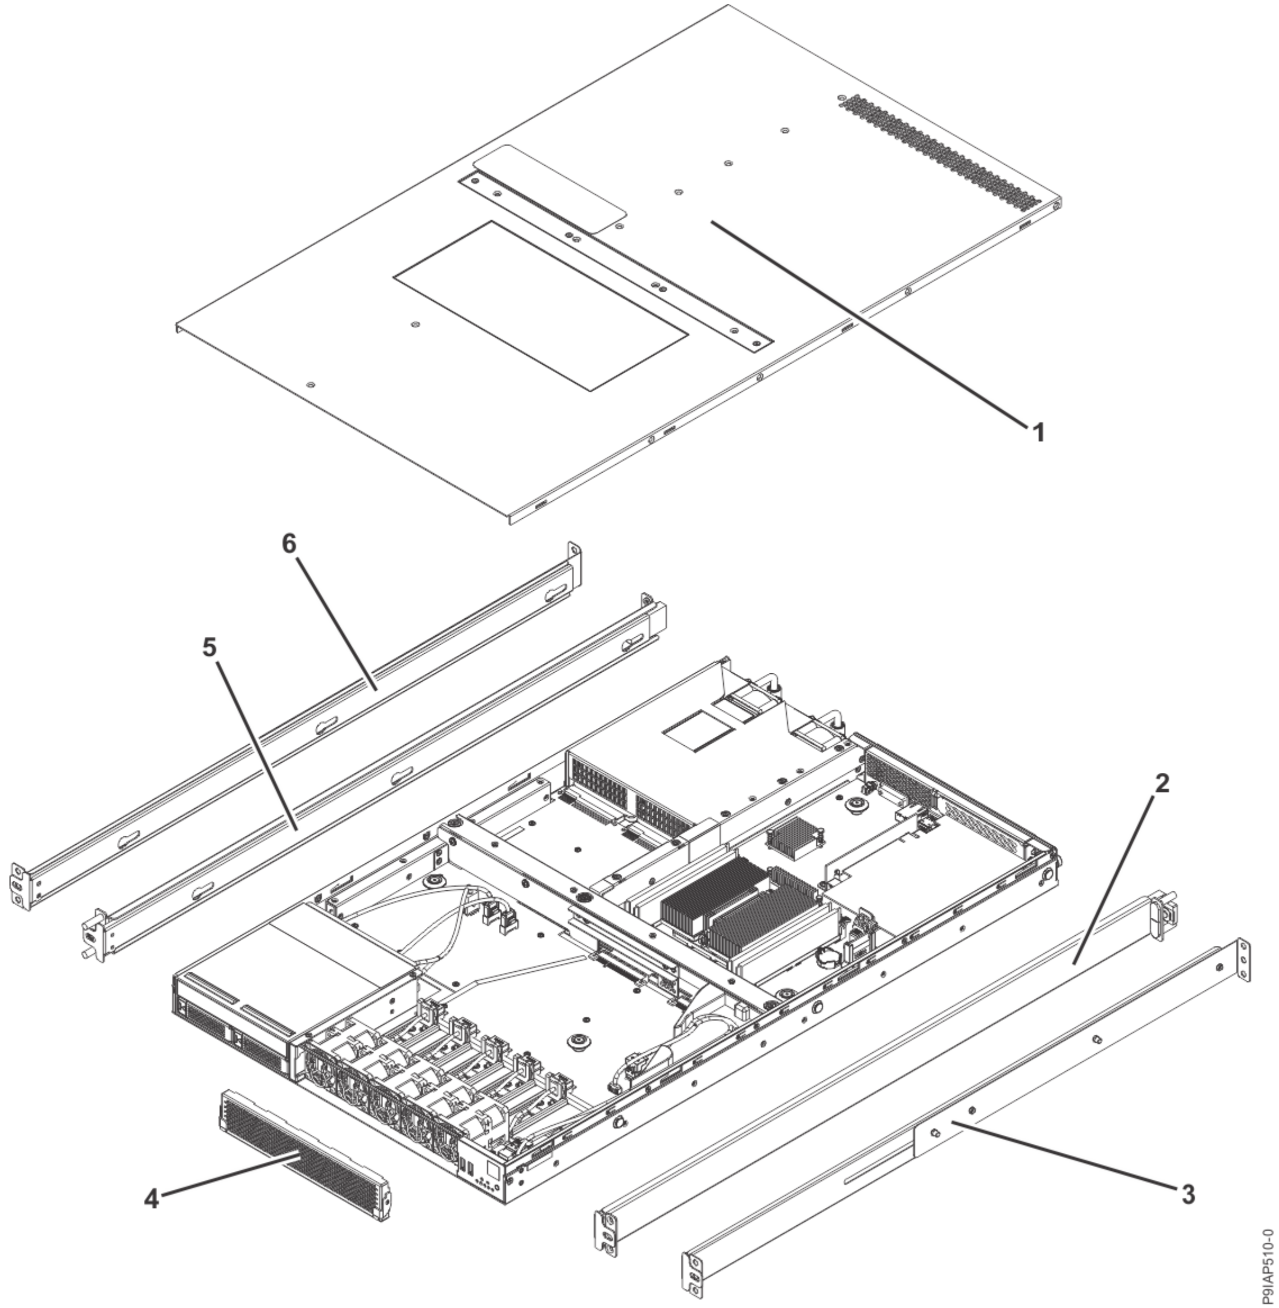

Figure 6. Rack final assembly

Table 10. Rack final assembly part numbers

Index number	Part number	Units per assembly	Description
1		1	Top cover assembly
2	03GM910	1	Fixed rail kit - contains left rail, right rail, attaching screws, and shipping brackets
3	03GM955	1	Fixed rail kit (adjustable in length) - contains left rail, right rail, and attaching screws
4	03GM764	1	Front bezel

Table 10. Rack final assembly part numbers (continued)

Index number	Part number	Units per assembly	Description
5	03GM910	1	Fixed rail kit - contains left rail, right rail, attaching screws, and shipping brackets
6	03GM955	1	Fixed rail kit (adjustable in length) - contains left rail, right rail, and attaching screws

System parts

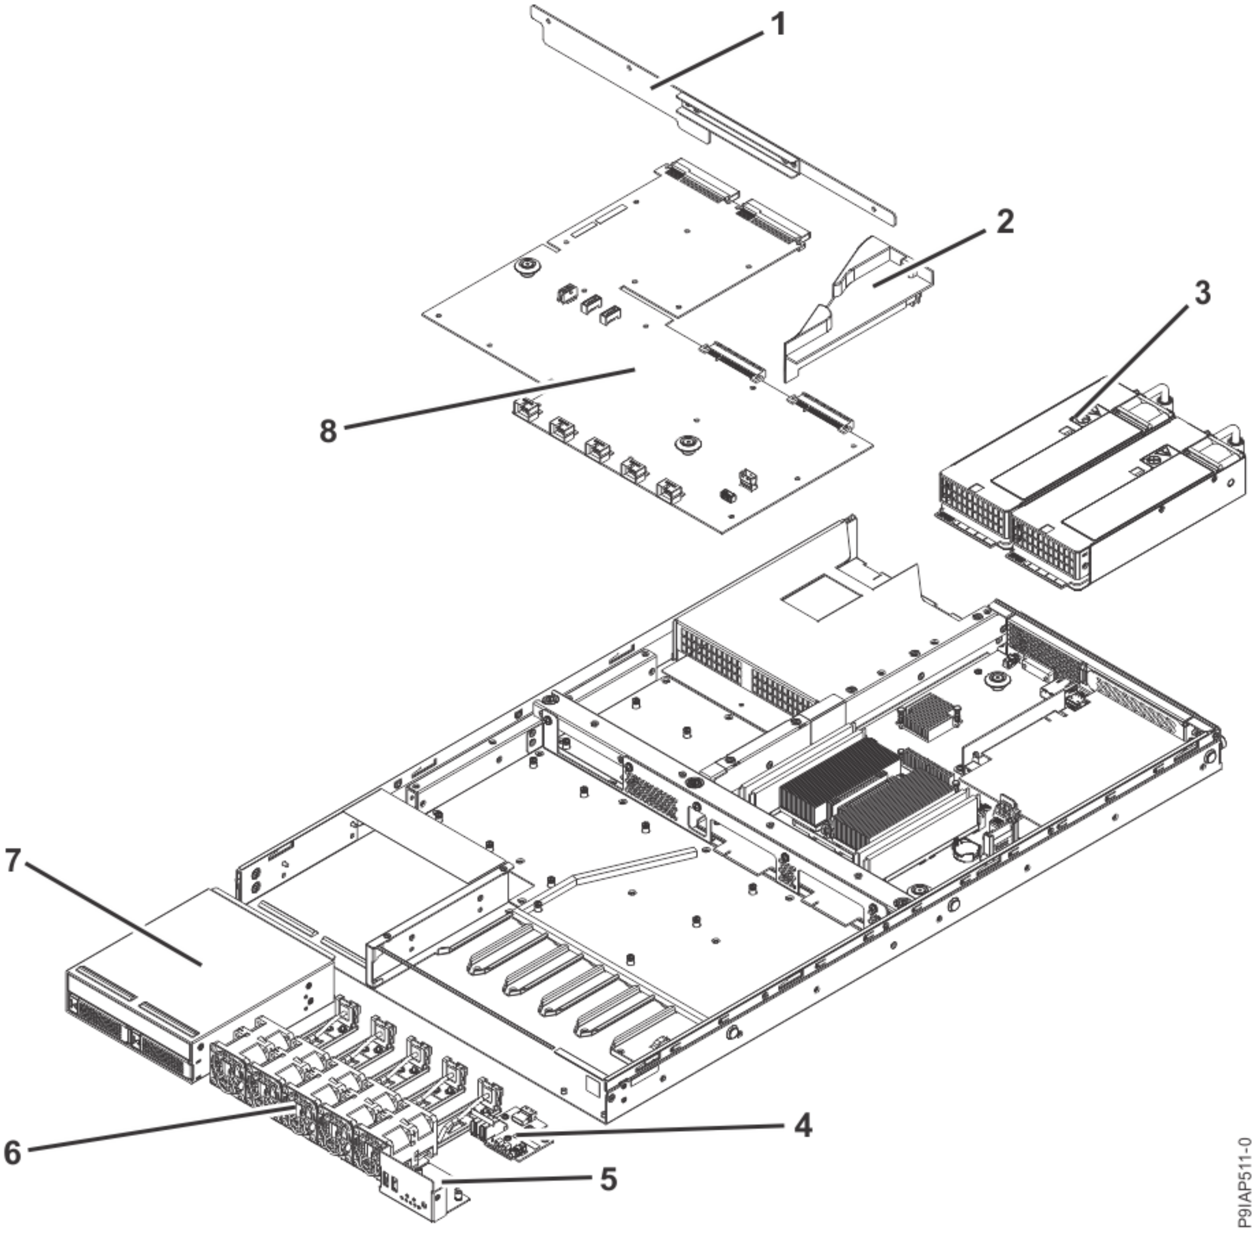

Figure 7. System parts

Table 11. System parts

Index number	Part number	Units per assembly	Description
1	03GM757	1	Air baffle (left)
2	03GM759	1	Air baffle (right)
3	03FP372	2	Power supply
4	02WF445	1	Control panel card
5		1	Control panel cover
6	03GM774	5	Fan
7	03GM776	1	Drive holder
	03GM792	2	1.8 TB 2.5 inch SAS disk drive (includes drive and carrier)
	03GM803	2	Drive carrier
8	02WF441	1	Power distribution board

Additional system parts

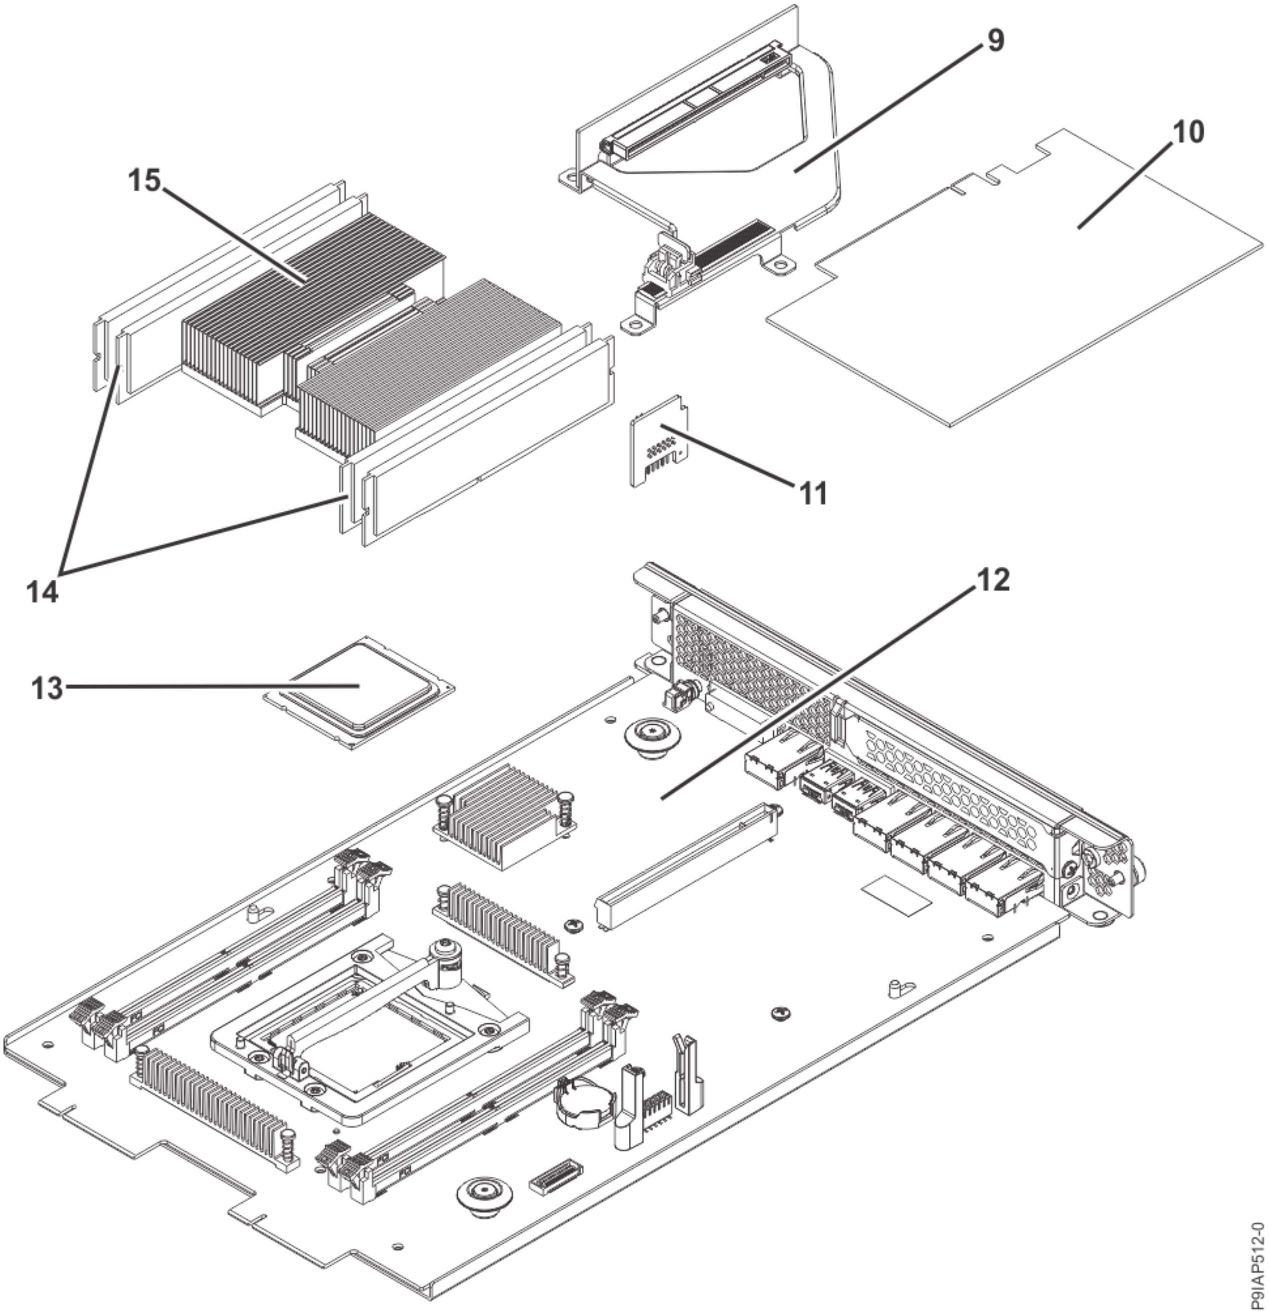

Figure 8. Additional system parts

Table 12. Additional system parts

Index numbe r	Part number	Units per assembly	Description
9	02WF443	1	PCIe riser
10	02JD569	1	PCIe2 2-port 10 GbE BaseT RJ45 adapter
11	00VK865	1	Trusted platform module

Table 12. Additional system parts (continued)

Index number	Part number	Units per assembly	Description
12	02WF439	1	System backplane kit (includes time-of-day battery, PCIe riser, PCIe tailstock filler, system processor module removal tool, and thermal interface material)
13	03GM808	1	6 core 3.0 GHz system processor module kit (includes system processor module, thermal interface material, tweezers, and system processor module removal tool)
14	78P6722	4	16 GB 2RX4 DDR4 IS RDIMM (Micron Technology, Inc.)
	78P4197	4	16 GB 2RX4 DDR4 IS RDIMM (Samsung Electronics Co., Ltd. or SK hynix, Inc.)
	78P4198	4	32 GB 2RX4 DDR4 IS RDIMM (Micron Technology, Inc., Samsung Electronics Co., Ltd., or SK hynix, Inc.)
15	03GM989	1	Heat sink kit (includes heat sink and thermal interface material)

Table 13. Miscellaneous parts

Description	Part number
USB cable	03GM778
Control panel cable	03GM780
Drive signal cable	03GM782
Drive power cable	03GM784
CR2032 Lithium time-of-day battery	00RY543
Chassis crossbar	03GM755

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with ICT Accessibility 508 Standards and 255 Guidelines (https://www.access-board.gov/ict/) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power servers.

The IBM Power servers online product documentation in IBM Documentation is enabled for accessibility. For more information about IBM's commitment to accessibility, see the IBM accessibility website at IBM Accessibility (https://www.ibm.com/able/).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Electronic emission notices

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Class A Notices

The following Class A statements apply to the IBM servers that contain the Power10 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

The following Class A statements apply to the servers.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

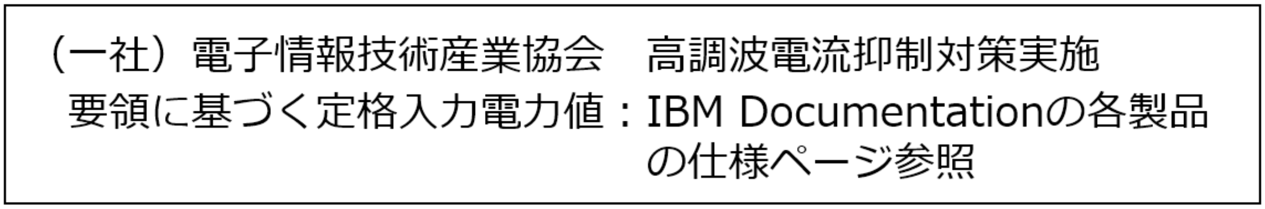

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：IBM Documentationの各製品
の仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

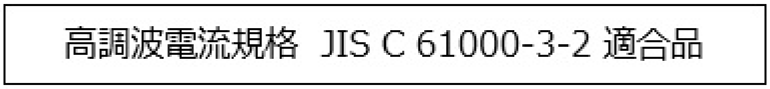

高調波電流規格　JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

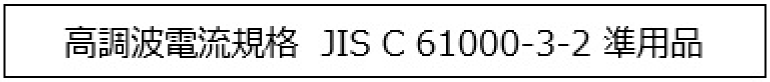

高調波電流規格　JIS C 61000-3-2 準用品

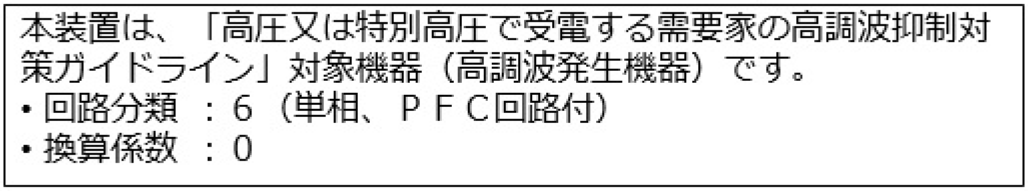

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

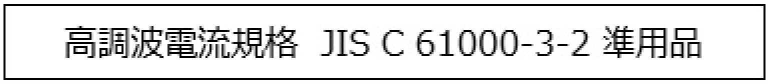

高調波電流規格　JIS C 61000-3-2 準用品

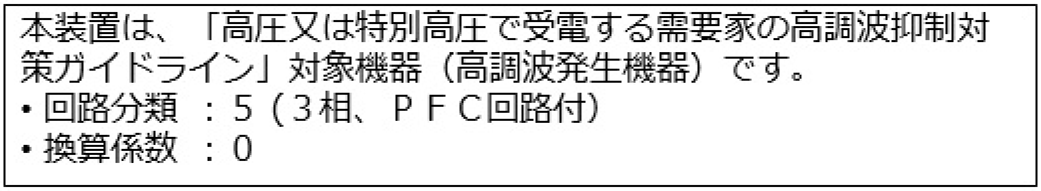

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

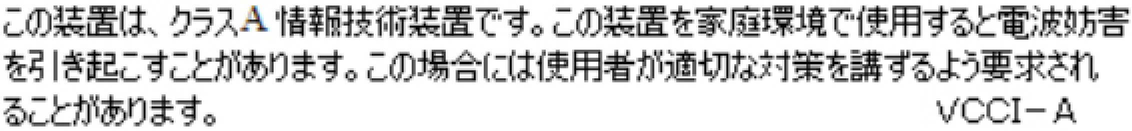

この装置は、クラスA 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。　VCCI－A

Korea Notice

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서
가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

警　告

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰
在这种情况下，可能需要用户对
其干扰采取切实可行的措施

Russia Notice

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

警告使用者：
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

United Kingdom Notice

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

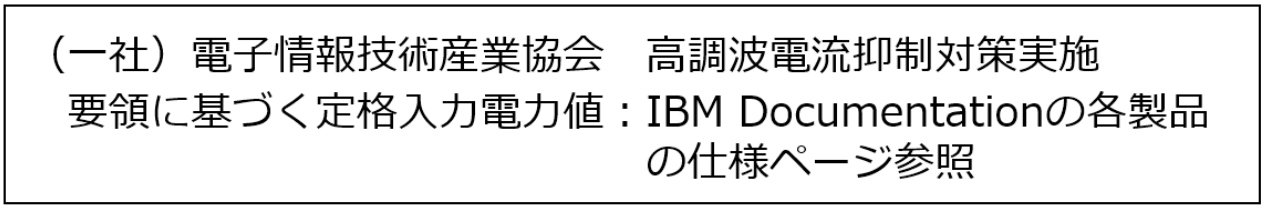

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：IBM Documentationの各製品の仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

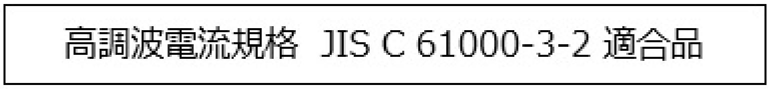

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

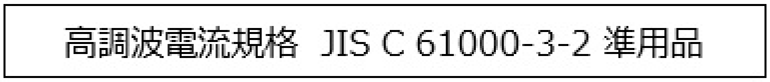

高調波電流規格 JIS C 61000-3-2 準用品

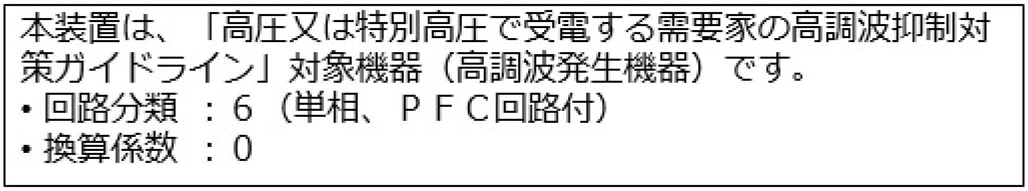

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：6（単相、ＰＦＣ回路付）
- 換算係数 ：0

This statement applies to products greater than 20 A per phase, three-phase.

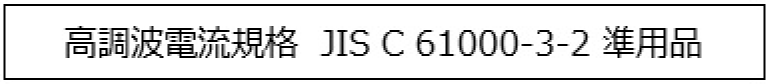

高調波電流規格 JIS C 61000-3-2 準用品

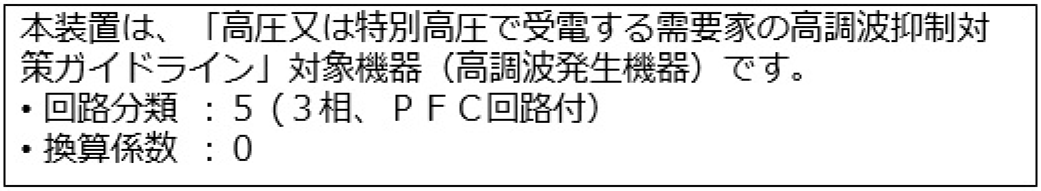

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 ：5（3相、ＰＦＣ回路付）
- 換算係数 ：0

Japan Voluntary Control Council for Interference (VCCI) Notice

この装置は，クラスB情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。　　　VCCI－B

Taiwan Notice

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Power Systems

Servicing the IBM Power Systems HMC (7063-CR2)

IBM

Note

Before using this information and the product it supports, read the information in "Safety notices" on page v, "Notices" on page 109, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125–5823.

This edition applies to IBM® Power® servers that contain the Power10 processor and to all associated models.

Contents

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

⚠ **DANGER:** When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.

- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.
- 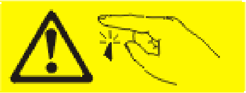 Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).

- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.

  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer's DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

⚠ **CAUTION:**

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

⚠ **CAUTION:** Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
  - Remove all devices in the 32U position (compliance ID RACK-001 or 22U (compliance ID RR001) and above.
  - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.

  - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U (compliance ID RACK-001 or 22U (compliance ID RR001) level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 2083 mm (30 x 82 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
  - Lower the four leveling pads.
  - Install stabilizer brackets on the rack cabinet or in an earthquake environment bolt the rack to the floor.
  - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)

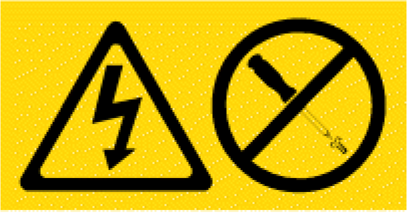

⚠ **DANGER:** Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)

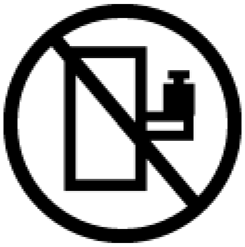

⚠ **DANGER:** Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack-mounted devices and do not use them to stabilize your body position (for example, when working from a ladder). Stability hazard:

- The rack may tip over causing serious personal injury.

- Before extending the rack to the installation position, read the installation instructions.
- Do not put any load on the slide-rail mounted equipment mounted in the installation position.
- Do not leave the slide-rail mounted equipment in the installation position.

(L002)

(L003)

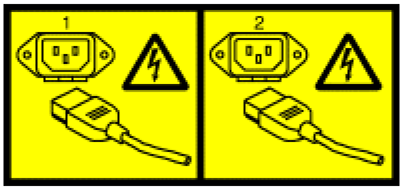

or

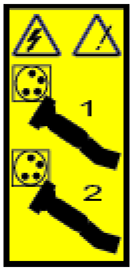

or

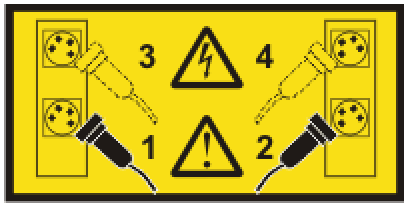

or

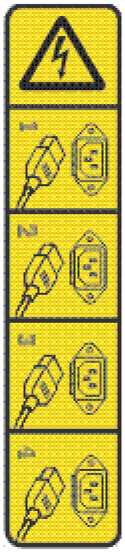

or

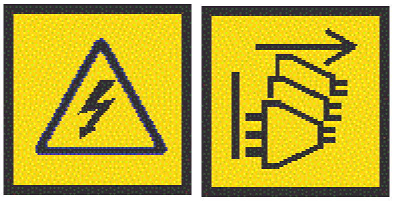

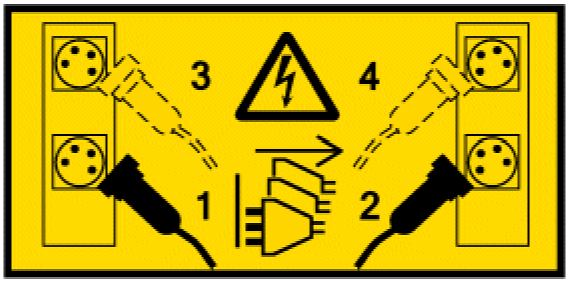

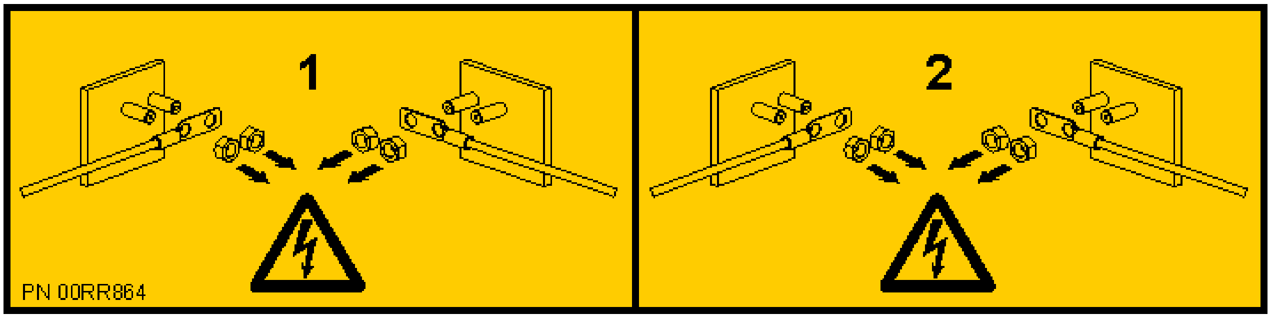

⚠ **DANGER:** Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)

⚠ **CAUTION:** A hot surface nearby. (L007)

(L008)

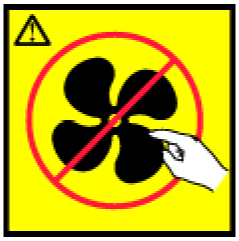

⚠ **CAUTION:** Hazardous moving parts nearby. (L008)

(L018)

or

⚠ **CAUTION:** High levels of acoustical noise are (or could be under certain circumstances) present. Use approved hearing protection and/ or provide mitigation or limit exposure. (L018)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

⚠ **CAUTION:** This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

⚠ **CAUTION:** Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

⚠ **CAUTION:** This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

⚠ **CAUTION:** Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information:

- Laser radiation when open.
- Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

(C030)

⚠ **CAUTION:** The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- Throw or immerse into water
- Heat to more than 100 degrees C (212 degrees F)
- Repair or disassemble

Exchange only with the IBM-approvedapproved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)(C003a)

⚠ **CAUTION:** Regarding IBM providedprovided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.

- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platforms, tilt riser, angled unit install wedge or other such accessory options. Secure such platforms -- riser tilt, wedge, etc options to main lift shelf or forks in all four (4x or all other provisioned mounting) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt [adjustable angling platform] option flat at all times except for final minor angle adjustment when needed.
- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL (unless the specific allowance is provided for one following qualified procedures for working at elevations with this TOOL).
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury.
- This TOOL must be maintained correctly for IBM Service personnel to use it. IBM shall inspect condition and verify maintenance history before operation. Personnel reserve the right not to use TOOL if inadequate. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intra-building ports of this equipment are suitable for connection to intra-building or unexposed wiring or cabling only. The intra-building ports of this equipment *must not* be metallically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metallically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The AC-powered system does not require the use of an external surge protection device (SPD).

The DC-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The DC-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

Removing and replacing parts in the 7063-CR2

Use these procedures to remove and replace failing parts in the IBM Power Systems HMC (7063-CR2). These parts are referred to as field replaceable units (FRUs).

Note: See the International Information Bulletin for Customers - Installation of IBM Machines (http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss). This bulletin (Publication number SC27-6601-00) provides a list of the key IBM system installation activities and those activities that might be billable activities.

Before you begin a replacement, complete these tasks:

1. If you are completing a replacement procedure that might put your data at risk, ensure that you have a current backup of your system or logical partition (including operating systems, licensed programs, and data).
2. Review the installation or replacement procedure for the feature or part.
3. Blue on a part of the hardware indicates a touch point where you can grip the hardware to remove it from or install it in the system, or open or close a latch.
4. Ensure that you have access to a medium, flat-blade screwdriver, and a Phillips screwdriver.
5. If parts are incorrect, missing, or visibly damaged, contact the provider of the part or your next level of support.

⚠ **DANGER:** When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5) For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.

- Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

⚠ **Attention:**

Failure to follow the step-by-step sequence for FRU removal or installation might result in FRU or system damage.

For safety, airflow purposes and thermal performance, the service access cover must be installed and fully seated before you power on the system.

For safety and airflow purposes and thermal performance, if you remove parts from the system, you must ensure that PCIe tail-stock fillers are present.

Use the following precautions whenever you handle electronic components or cables.

- The electrostatic discharge (ESD) kit and the ESD wrist strap must be used when you handle logic cards, single chip modules (SCM), multi-chip modules (MCM), electronic boards, and drives.
- Keep all electronic components in the shipping container or envelope until you are ready to install them.
- If you remove and then reinstall an electronic component, temporarily place the component on an ESD pad or blanket.

Removing and replacing cables in the 7063-CR2

Learn how to remove and replace cables in the IBM Power Systems HMC (7063-CR2) system.

Removing and replacing the drive power cable in the 7063-CR2

Learn how to remove and replace the drive power cable in the IBM Power Systems HMC (7063-CR2) system.

Removing the drive power cable from the 7063-CR2 system

To remove the drive power cable from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.

- When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
- If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

2. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.

3. Place the system in the service position on an ESD surface on a table.

   For instructions, see "Placing a 7063-CR2 system into the service position" on page 101.

4. Remove the service access cover.

   For instructions, see "Removing the service access cover from a 7063-CR2 system" on page 104.

5. Remove the drive power cable.

   a) Label where the drive power cable **(A)** connects to the power distribution board and to the drive holder.

      This cable is a "Y" cable that connects to the drive holder in two places.

   b) Remove the drive power cable from the power distribution board and from the drive holder.

      Use your thumb or finger to press the release latch on the connector to remove the cable from the power distribution board.

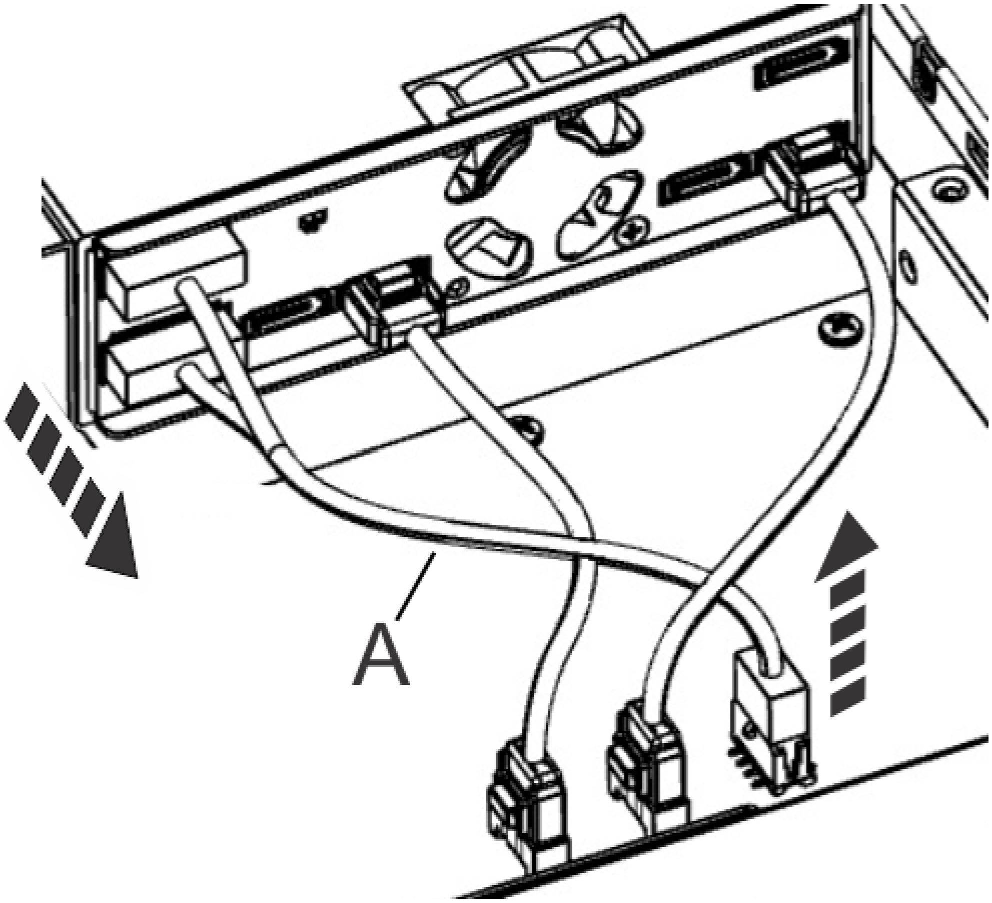

Figure 1. Disconnecting the drive power cable

Replacing the drive power cable in the 7063-CR2 system

To replace the drive power cable in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Using your labels, replace the drive power cable into the power distribution board and into the drive holder.

   Ensure that the cable latch clip snaps into place in the connector on the power distribution board.

   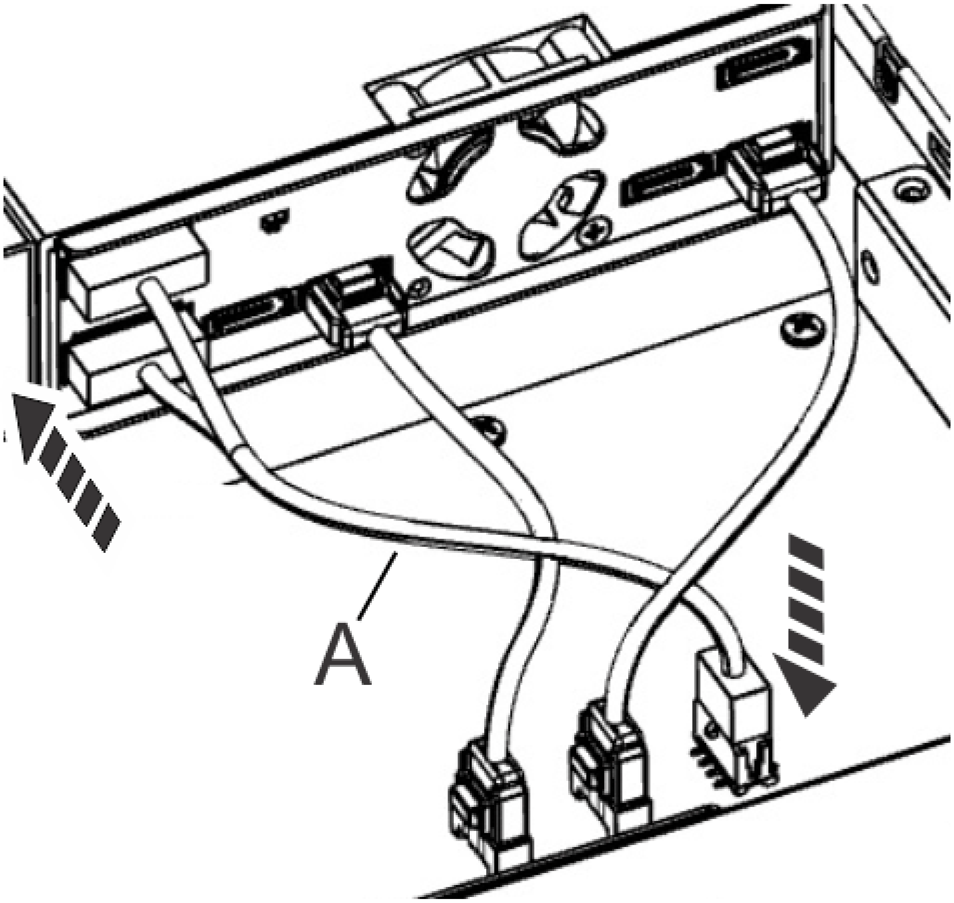

   Figure 2. Connecting the drive power cable
3. Install the service access cover.

   For instructions, see “Installing the service access cover on a 7063-CR2 system” on page 105.
4. Replace the system in the rack and replace the components that you removed.

   For instructions, see “Placing a 7063-CR2 system into the operating position” on page 103.
5. Power on the system for operation.

   For instructions, see “Starting the 7063-CR2 system ” on page 98.

Removing and replacing a drive signal cable in the 7063-CR2

Learn how to remove and replace a drive signal cable in the IBM Power Systems HMC (7063-CR2) system.

Removing a drive signal cable from the 7063-CR2 system

To remove a drive signal cable from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   Attention:

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.
2. The system has two drive signal cables. Identify which drive cable has the issue.
3. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.
4. Place the system in the service position on an ESD surface on a table.

   For instructions, see "Placing a 7063-CR2 system into the service position" on page 101.
5. Remove the service access cover.

   For instructions, see "Removing the service access cover from a 7063-CR2 system" on page 104.
6. Remove the previously identified drive signal cable.

   a) Label where the drive signal cable **(A)** connects to the power distribution board and to the drive holder.

   b) Remove the drive signal cable from the power distribution board and from the drive holder.

   Use your thumb or finger to press the release latch on the connector to remove the cable.

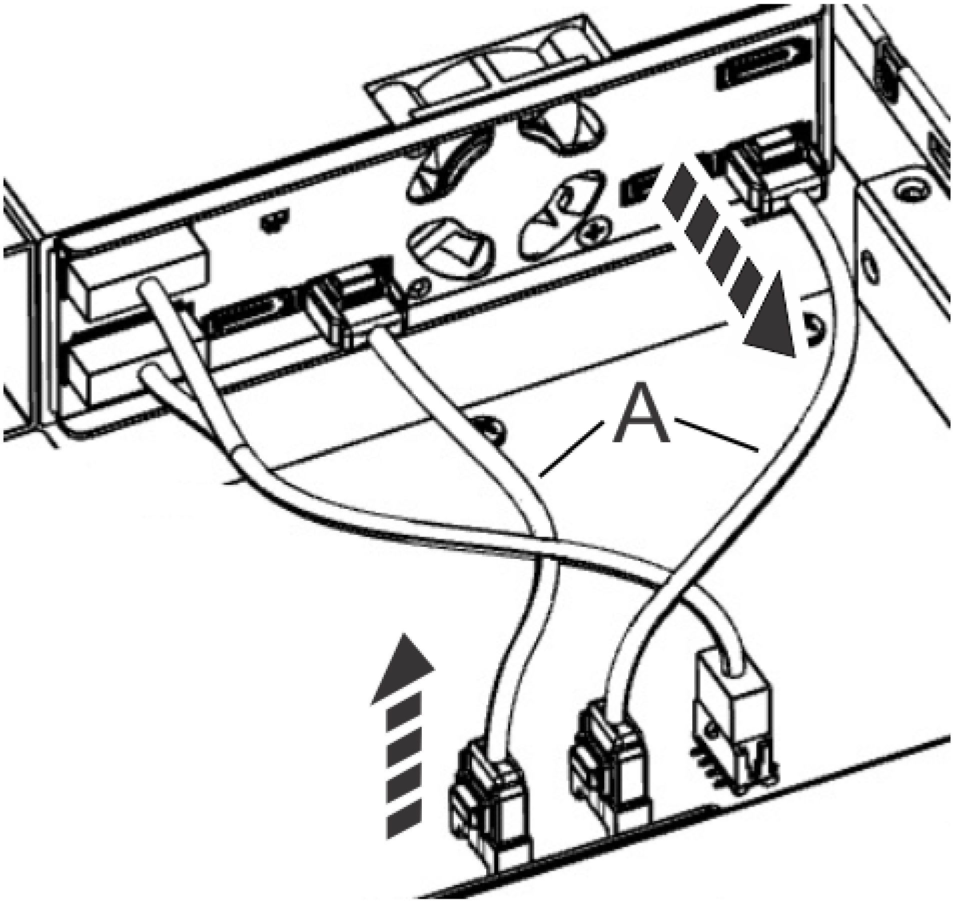

Figure 3. Disconnecting the drive signal cable

Replacing a drive signal cable in the 7063-CR2 system

To replace a drive signal cable in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Using your labels, replace the drive signal cable into the power distribution board and into the drive holder.

   Ensure that the cable latch clip snaps into place on the connector.

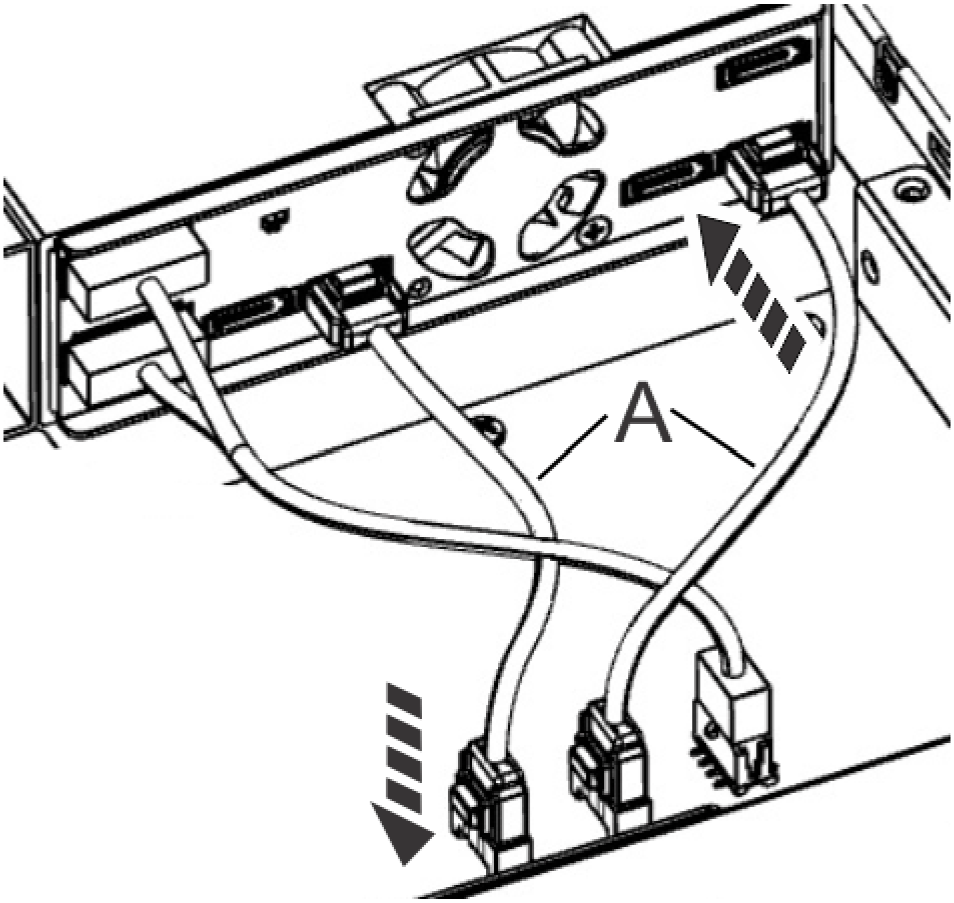

Figure 4. Connecting the drive signal cable

3. Install the service access cover.

   For instructions, see "Installing the service access cover on a 7063-CR2 system" on page 105.
4. Replace the system in the rack and replace the components that you removed.

   For instructions, see "Placing a 7063-CR2 system into the operating position" on page 103.
5. Power on the system for operation.

   For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the control panel cable in the 7063-CR2

Learn how to remove and replace the control panel cable in the IBM Power Systems HMC (7063-CR2) system.

Removing the control panel cable from the 7063-CR2 system

To remove the control panel cable from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.

- When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
- If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

2. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.
3. Place the system in the service position on an ESD surface on a table.

   For instructions, see "Placing a 7063-CR2 system into the service position" on page 101.
4. Remove the service access cover.

   For instructions, see "Removing the service access cover from a 7063-CR2 system" on page 104.
5. Remove the control panel cable.
   a) Label where the control panel cable **(A)** connects to the power distribution board and to the control panel.
   b) Remove the control panel cable from the power distribution board and from the control panel.

      Use your thumb or finger to press the release latch on the connector to remove the cable. Unclip the cable from the right air baffle.

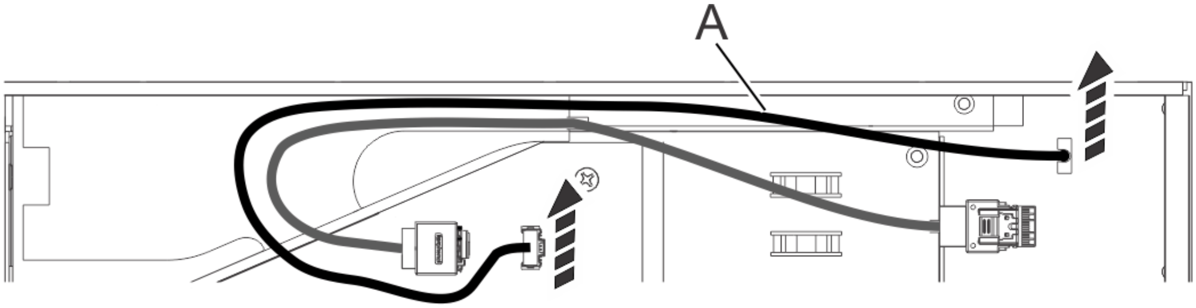

Figure 5. Disconnecting the control panel cable

Replacing the control panel cable in the 7063-CR2 system

To replace the control panel cable in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Replace the control panel cable.
   a) Using your labels, replace the control panel cable **(A)** into the power distribution board and into the control panel.

      Ensure that the cable latch clip snaps into place on the connector.
   b) Route the cable through the right air baffle channel; the USB cable nests inside the control panel cable.

      Clip the cable into the right air baffle.

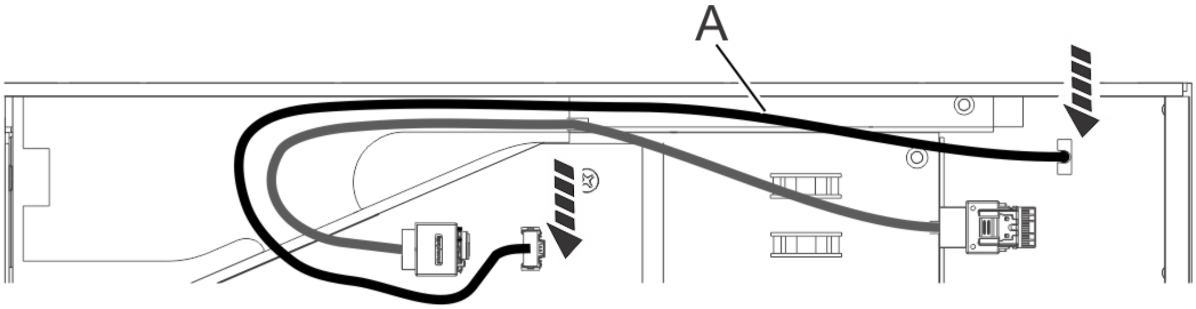

Figure 6. Connecting the control panel cable

3. Install the service access cover.

   For instructions, see "Installing the service access cover on a 7063-CR2 system" on page 105.
4. Replace the system in the rack and replace the components that you removed.

   For instructions, see "Placing a 7063-CR2 system into the operating position" on page 103.
5. Power on the system for operation.

   For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the USB cable in the 7063-CR2

Learn how to remove and replace the USB cable in the IBM Power Systems HMC (7063-CR2) system.

Removing the USB cable from the 7063-CR2 system

To remove the USB cable from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.
2. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.
3. Place the system in the service position on an ESD surface on a table.

   For instructions, see "Placing a 7063-CR2 system into the service position" on page 101.
4. Remove the service access cover.

   For instructions, see "Removing the service access cover from a 7063-CR2 system" on page 104.
5. Remove the USB cable.

   a) Label where the USB cable **(A)** connects to the power distribution board and to the control panel.

b) Remove the USB cable from the power distribution board and from the control panel.

Use your thumb or finger to press the release latch on the connector to remove the cable. Unclip the cable from the right air baffle.

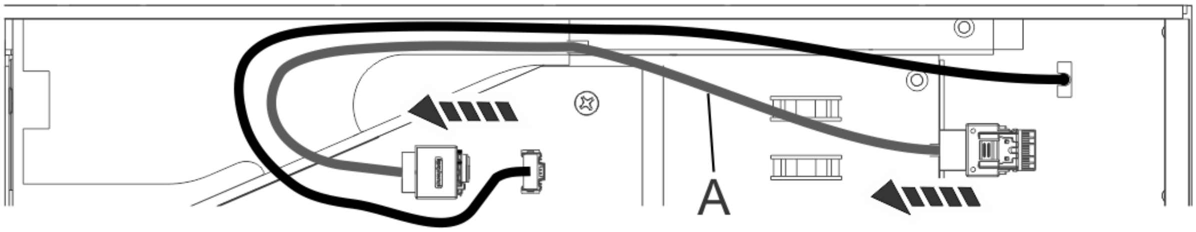

Figure 7. Disconnecting the USB cable

Replacing the USB cable in the 7063-CR2 system

To replace the USB cable in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Replace the USB cable.

   a) Using your labels, replace the USB cable into the power distribution board and into the control panel.

   Ensure that the cable latch clip snaps into place on the connector.

   b) Route the cable through the right air baffle channel; the USB cable nests inside the control panel cable.

   Clip the cable into the right air baffle.

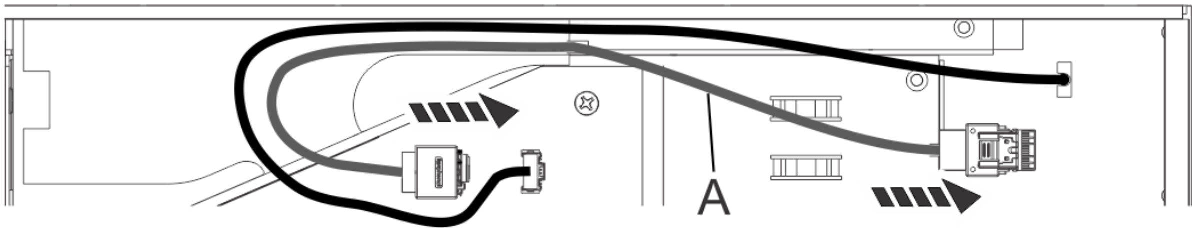

Figure 8. Connecting the USB cable

3. Install the service access cover.

   For instructions, see "Installing the service access cover on a 7063-CR2 system" on page 105.
4. Replace the system in the rack and replace the components that you removed.

   For instructions, see "Placing a 7063-CR2 system into the operating position" on page 103.
5. Power on the system for operation.

   For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing a drive in the 7063-CR2

Learn how to remove and replace a drive in the IBM Power Systems HMC (7063-CR2) system.

About this task

The system has two physical drives. These drives are configured as a single virtual drive, as a RAID1 array. The drives have an ID of 0. If one of the drives needs to be replaced, use the "Drive commands for the 7063-CR2 system" on page 99 to check the status of the drives, the IDs of the drives, and to rebuild the RAID1 array. The drive can be replaced with power on; the HMC continues to function normally.

Removing a drive from the 7063-CR2 system

To remove a drive from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   Attention:

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.
2. Remove the drive. Use the position information indicated by the service log.
   a) Push in the left side of the handle release latch **(A)** to unlock the drive bay handle **(B)**.
   b) Pull out the drive bay handle **(B)** toward you. If the drive bay handle is not all the way out, the drive cannot slide out of the system.
   c) Support the bottom of the drive as you slide it out of the system. Do not hold the drive by the handle.
   d) Place the drive on an ESD surface.

   Figure 9. Removing a drive

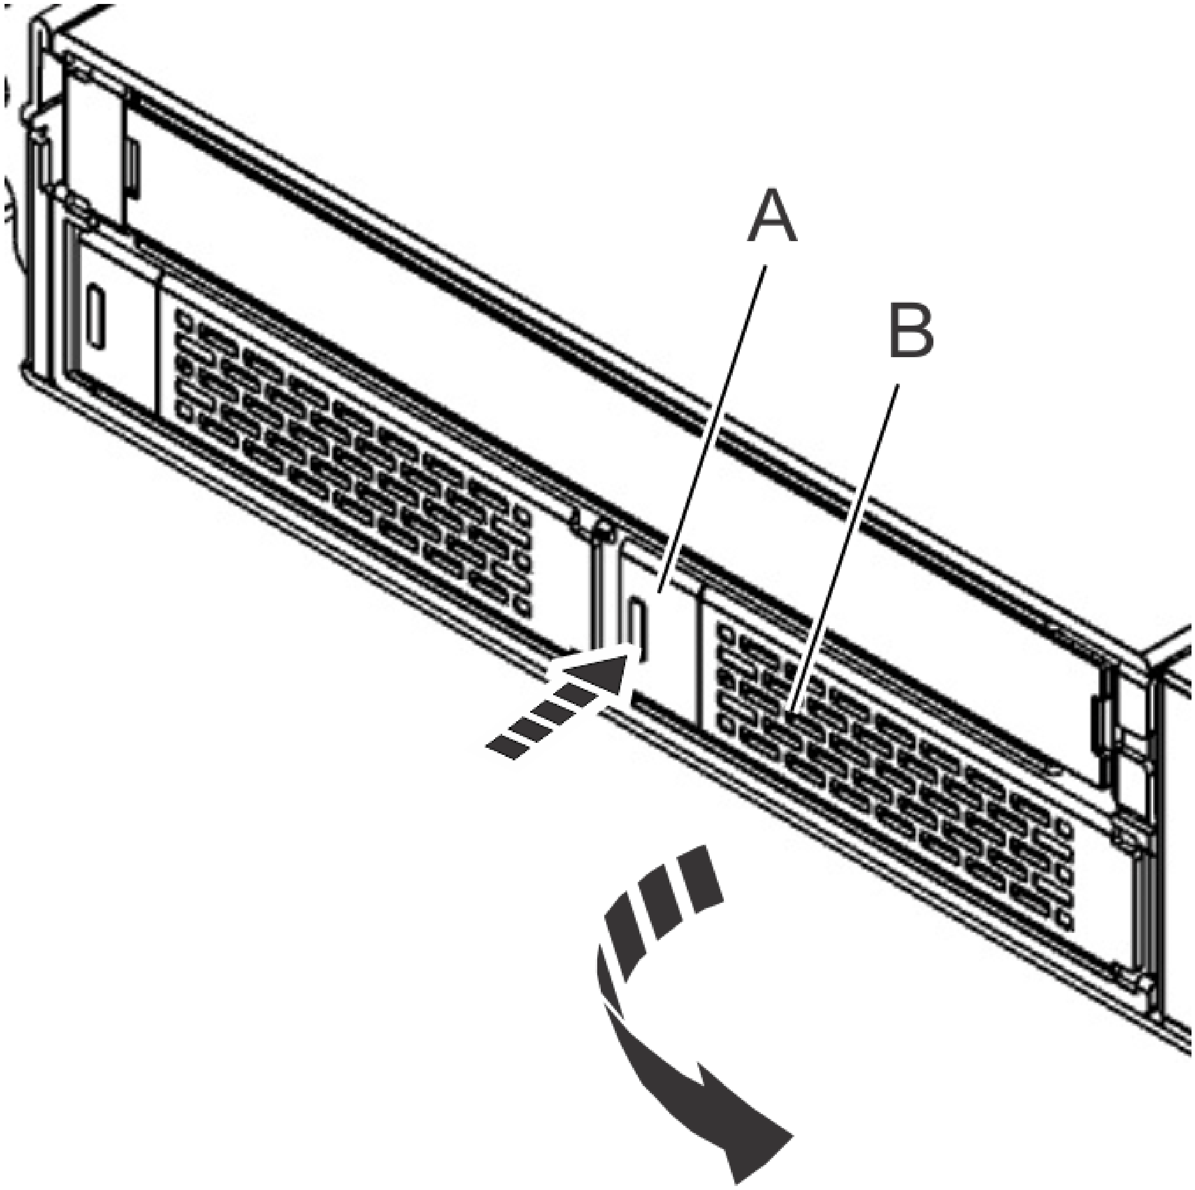

Replacing a drive in the 7063-CR2 system

To replace a drive in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Replace the drive.
   a) Support the drive by the bottom as you position the drive, and insert it into the drive slot.

      Important: Ensure that the drive is fully seated and is all the way into the system.

   b) Lock the drive bay handle **(A)** by pushing in the handle release latch until it locks into place.

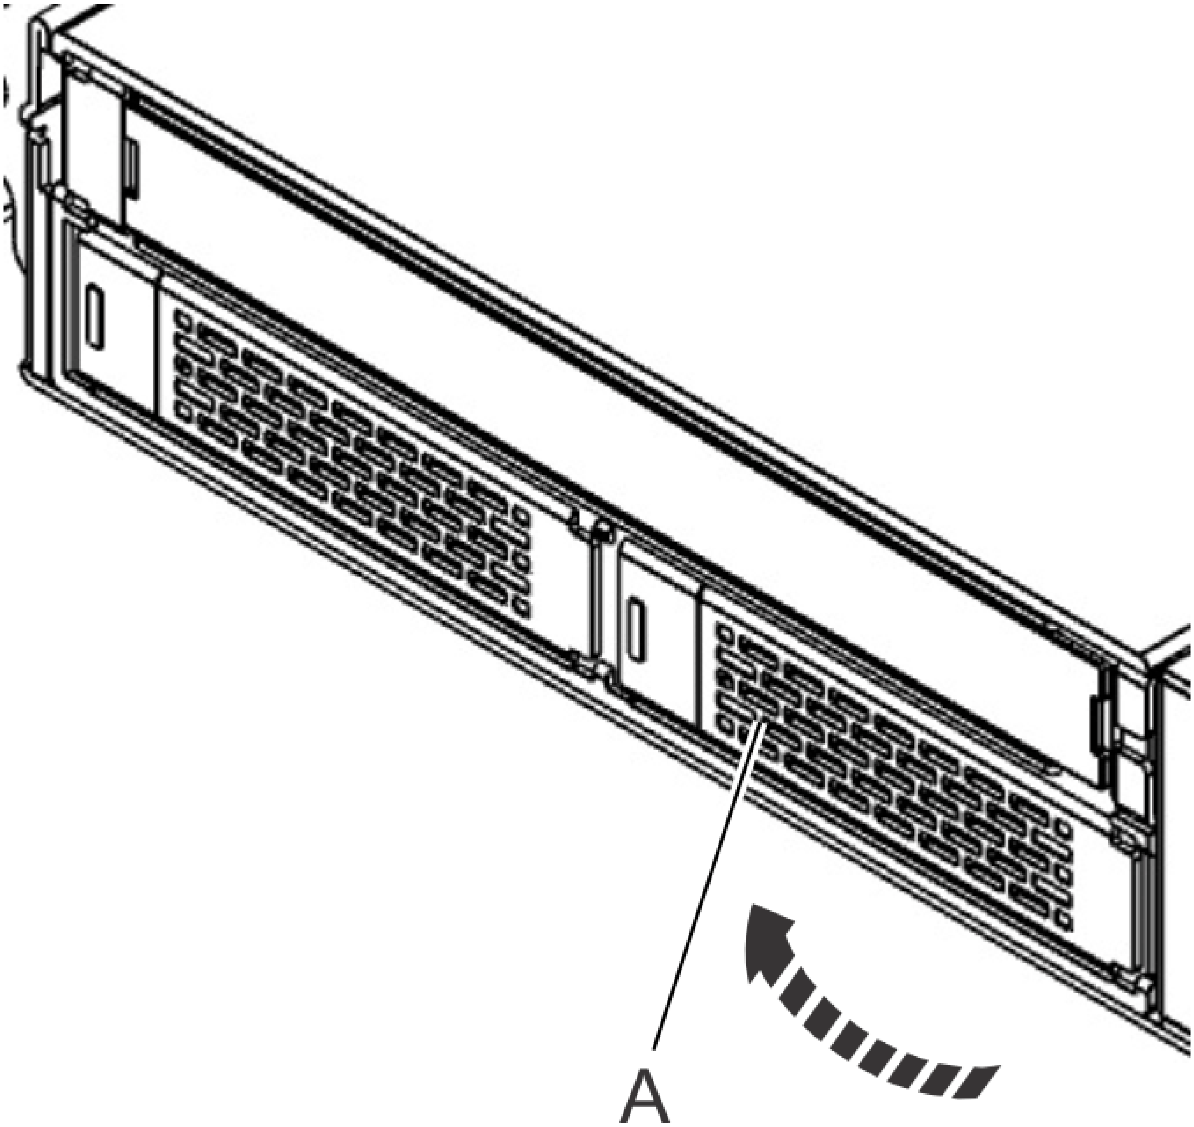

Figure 10. Installing a drive

3. The drive rebuild operation starts automatically.

Removing and replacing the drive holder in the 7063-CR2

Learn how to remove and replace the drive holder in the IBM Power Systems HMC (7063-CR2) system.

Removing the drive holder from the 7063-CR2 system

To remove the drive holder from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

You can use a magnetic tip screwdriver to remove and replace the screws.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

Attention:

- Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
- When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
- If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

2. Power off the system.

   For instructions, see “Stopping the 7063-CR2 system” on page 99.

3. Place the system in the service position on an ESD surface on a table.

   For instructions, see “Placing a 7063-CR2 system into the service position” on page 101.

4. Label and remove the two drives from the system.

   a) Push in the left side of the handle release latch **(A)** to unlock the drive bay handle **(B)**.

   b) Pull out the drive bay handle **(B)** toward you. If the drive bay handle is not all the way out, the drive cannot slide out of the system.

   c) Support the bottom of the drive as you slide it out of the system. Do not hold the drive by the handle.

   d) Place the drive on an ESD surface.

   Figure 11. Removing a drive

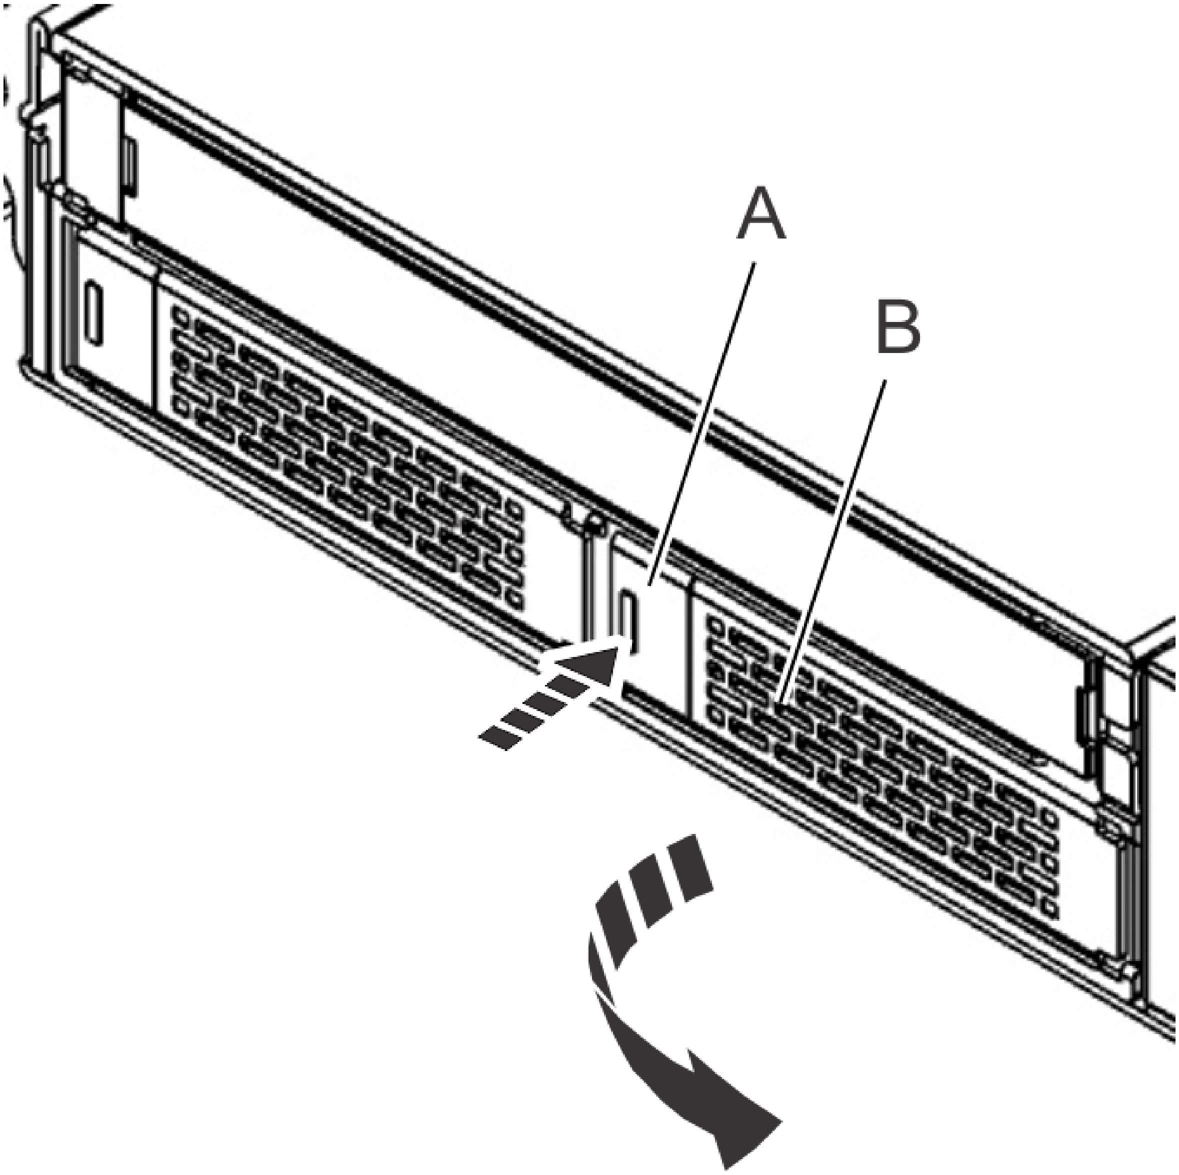

5. Remove all of the fans from the system.

   For instructions, see "Removing a fan from the 7063-CR2 system" on page 18.

6. Remove the service access cover.

   For instructions, see "Removing the service access cover from a 7063-CR2 system" on page 104.

7. Label and disconnect the cables from the rear of drive holder.

   For instructions, see "Removing the drive power cable from the 7063-CR2 system" on page 2 and "Removing a drive signal cable from the 7063-CR2 system" on page 5.

8. Remove left rail from the system.
9. Remove the four screws from the drive holder. The holder has two screws on each side.

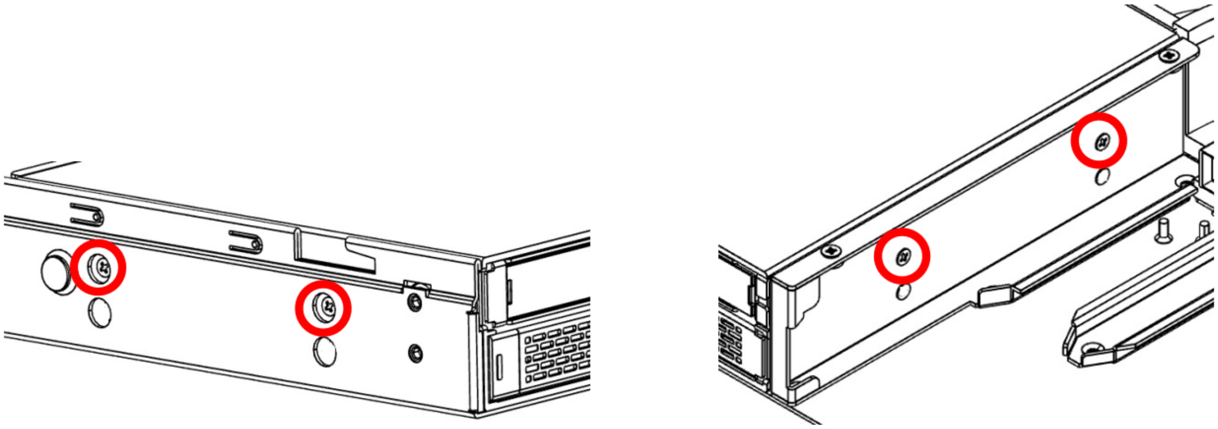

Figure 12. Drive holder screws

10. Remove the drive holder from the system.
11. Place the drive holder and cables on the table.

Replacing the drive holder in the 7063-CR2 system

To replace the drive holder in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

You can use a magnetic tip screwdriver to remove and replace the screws.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Insert the drive holder into the system.
3. Install the four screws for the drive holder. The holder has two screws on each side; use the top screw holes.

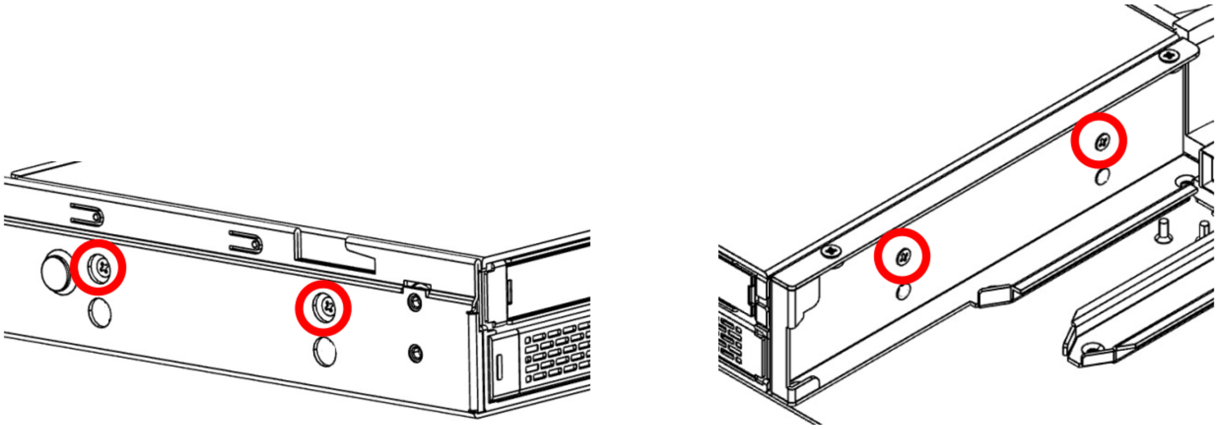

Figure 13. Drive holder screws

4. Install the left rail onto the system.
5. Using your labels, replace the cables into the rear of the drive holder.
6. Install the service access cover.

   For instructions, see "Installing the service access cover on a 7063-CR2 system" on page 105.
7. Replace all of the fans into the system.

For instructions, see “Replacing a fan in the 7063-CR2 system” on page 20.

8. Using your labels, replace the two drives into the system.
   a) Support the drive by the bottom as you position the drive, and insert it into the drive slot.

      Important: Ensure that the drive is fully seated and is all the way into the system.

   b) Lock the drive bay handle **(A)** by pushing in the handle release latch until it locks into place.

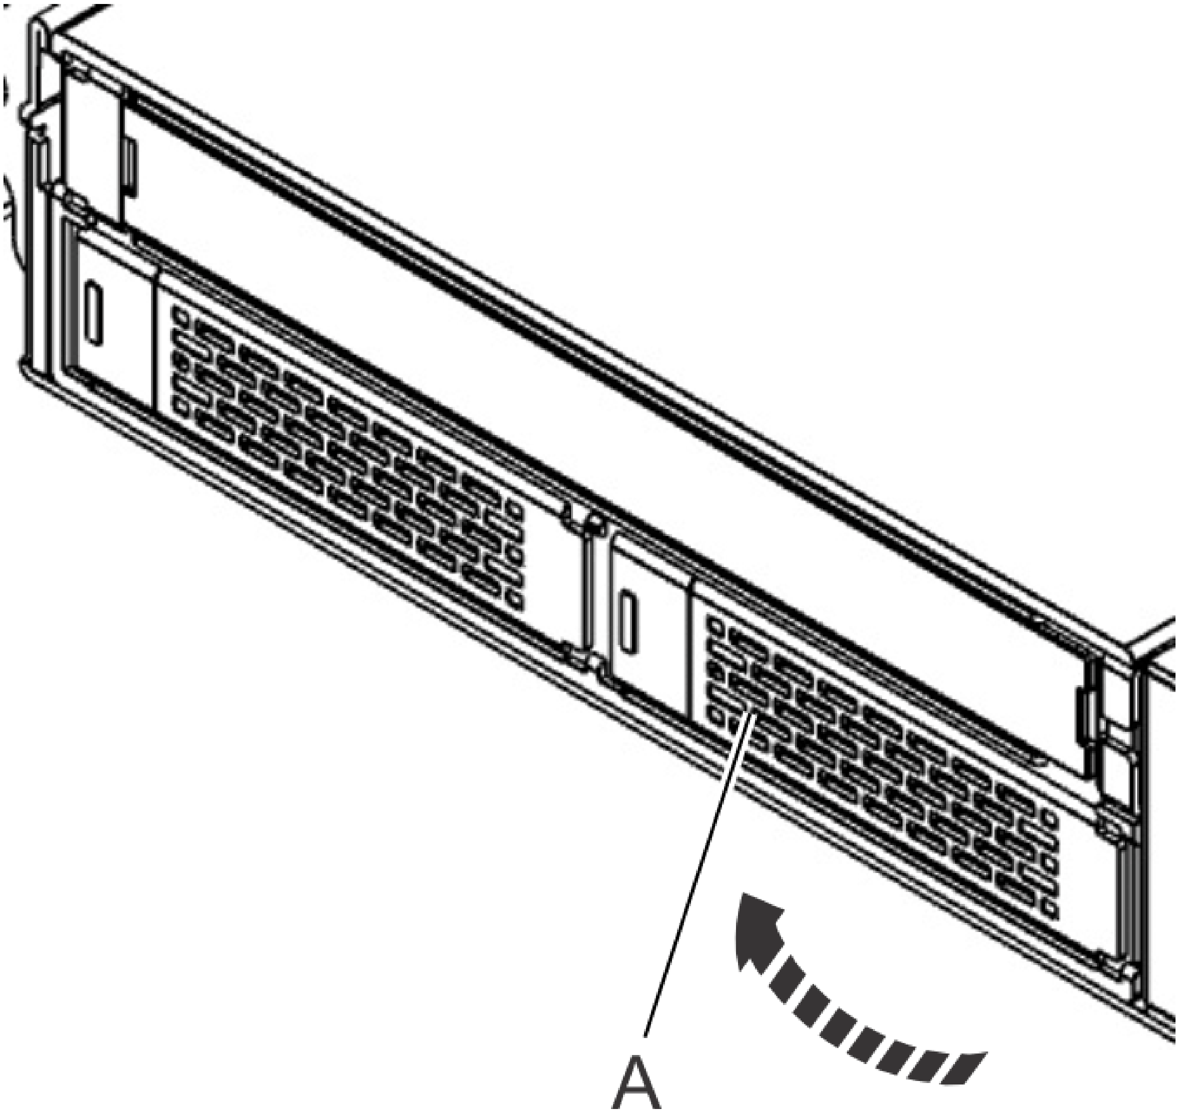

Figure 14. Installing a drive

9. Replace the system in the rack and replace the components that you removed.

   For instructions, see “Placing a 7063-CR2 system into the operating position” on page 103.

10. Power on the system for operation.

    For instructions, see “Starting the 7063-CR2 system ” on page 98.

Removing and replacing fans in the 7063-CR2

Learn how to remove and replace fans in the IBM Power Systems HMC (7063-CR2) system.

Removing a fan from the 7063-CR2 system

To remove a fan from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Before you begin

(L008)

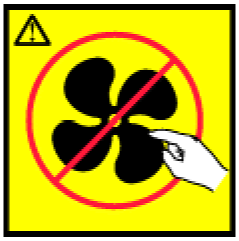

⚠ **CAUTION:** Hazardous moving parts nearby. (L008)

About this task

If a single fan failed, it can be replaced while the system is running.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.
2. Remove the fan cover from the front of the system.

   a) Rotate the two levers **(A)** on each side of the fan cover up and out to unlock the fan cover.

   b) Pull the cover away from the system.

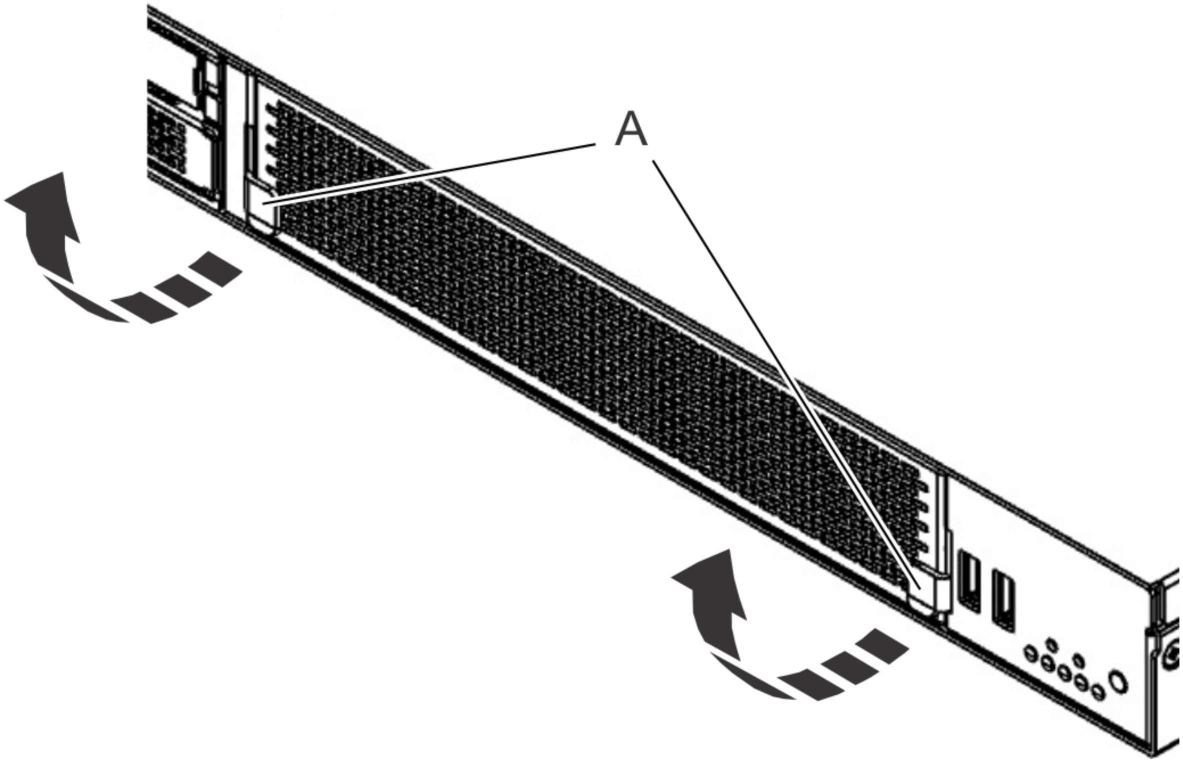

Figure 15. Removing the fan cover

3. Remove the fan assembly from the system.

   ⚠ **Warning:** If the system is powered on and you remove two or more fans, the system powers down.

   a) Use the ring **(A)** on the front of the fan to pull the fan from the system.

   b) Support the bottom of the fan as you slide it out of the system. Do not hold the fan by the ring.

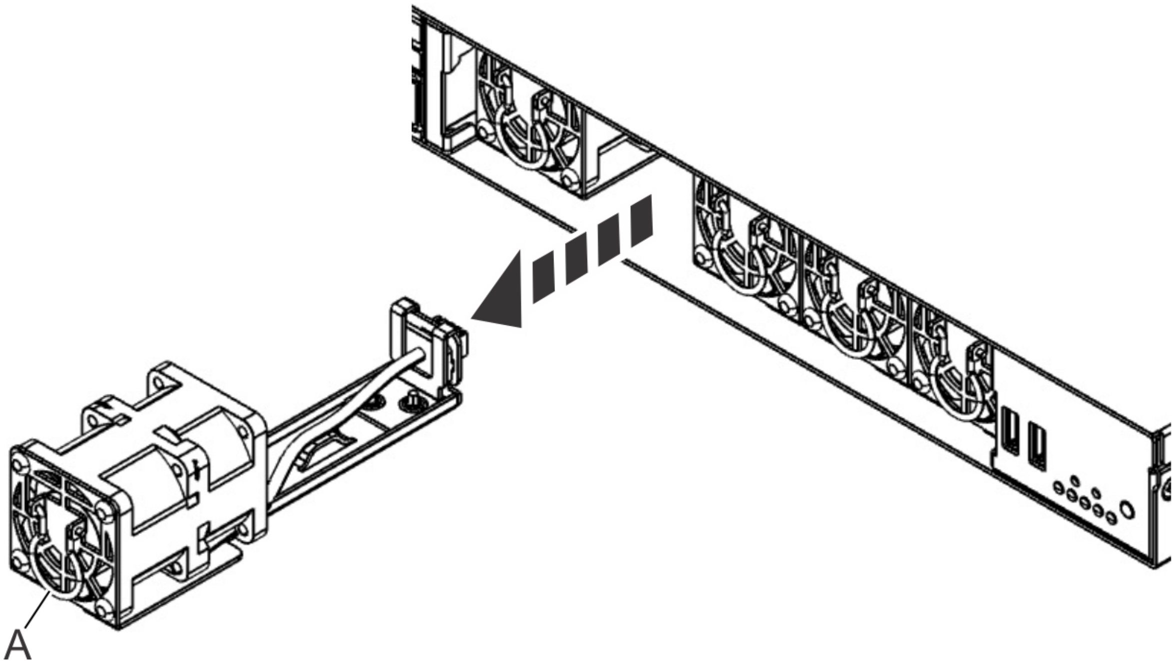

Figure 16. Removing a fan

Replacing a fan in the 7063-CR2 system

To replace a fan in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Replace the fan in the system. The top side of the fan assembly has an arrow that shows the airflow direction; the metal alignment plate is the bottom of the fan assembly. Ensure that you align the fan housing under the internal fan rails. Push the fan into the system until the fan is fully seated.

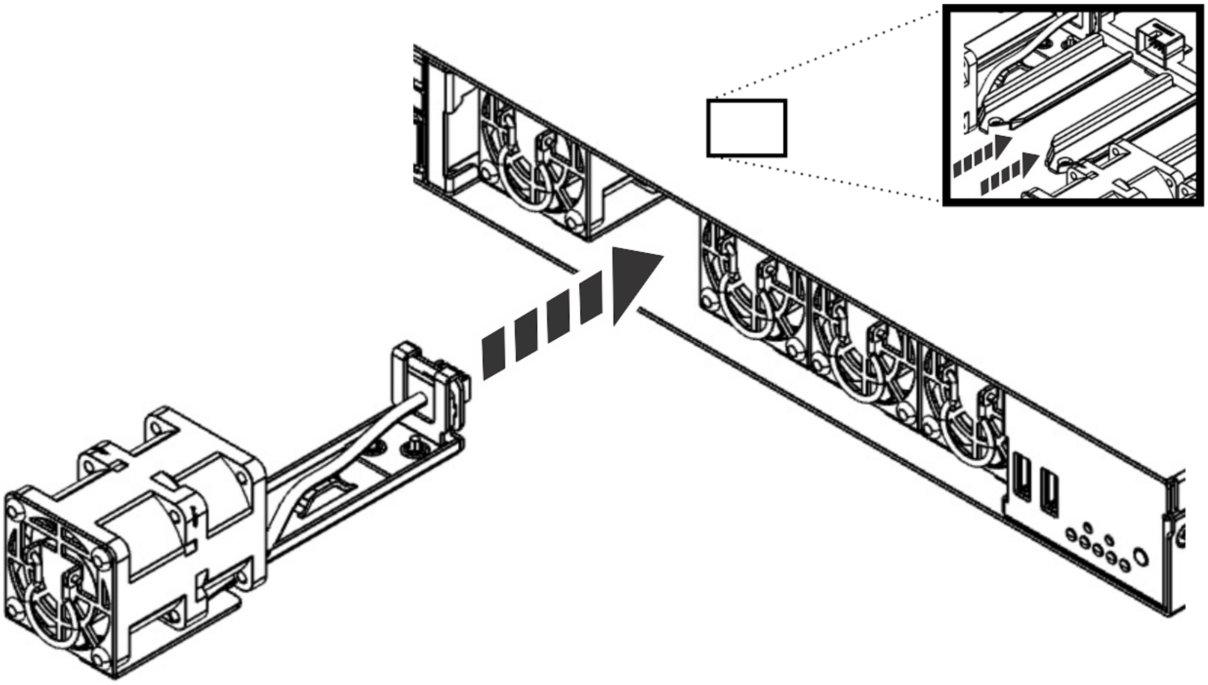

Figure 17. Replacing a fan

3. Replace the fan cover.
   a) Ensure that the two levers are open.
   b) Put the fan cover in place.
   c) Rotate the two levers on each side of the fan cover down and in to secure the cover to the system.

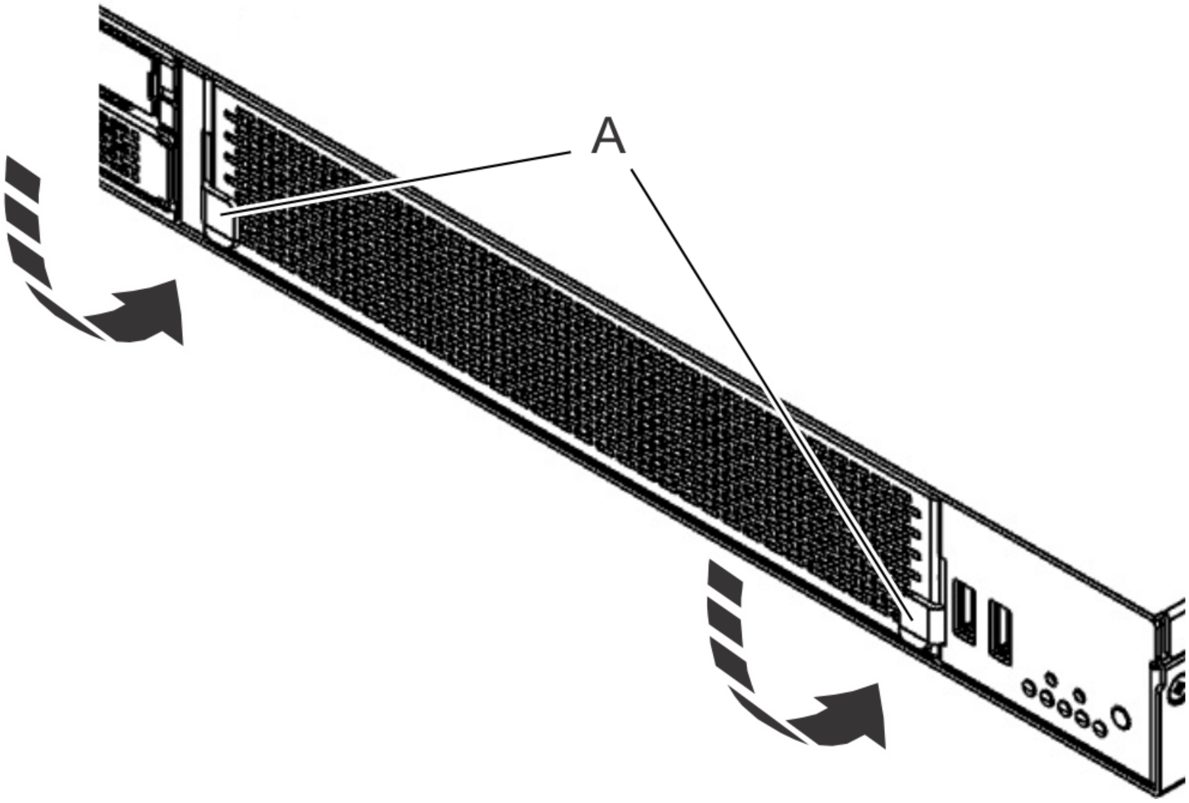

Figure 18. Replacing the fan cover

Removing and replacing memory in the 7063-CR2

To remove and replace memory in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

The four memory modules must be the same size and type. Mixing of memory module types is not allowed. The following table lists the supported memory feature codes.

Table 1. Memory feature codes

Feature code	Size
EM62	4 x 16 GB = 64 GB
EM63	4 x 32 GB = 128 GB

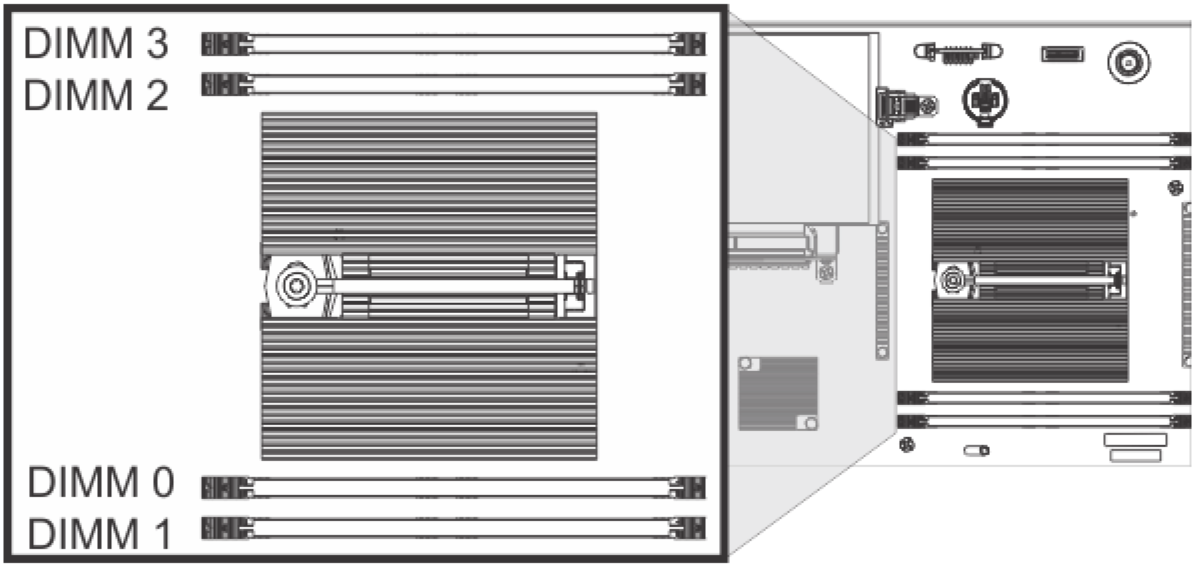

Figure 19. Memory locations

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.

2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   Attention:

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

3. Remove the system backplane from the rear of the system.

   a) Label and remove the two power cables.

   For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.

   b) Label and remove the signal cables from the rear of the system.

   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

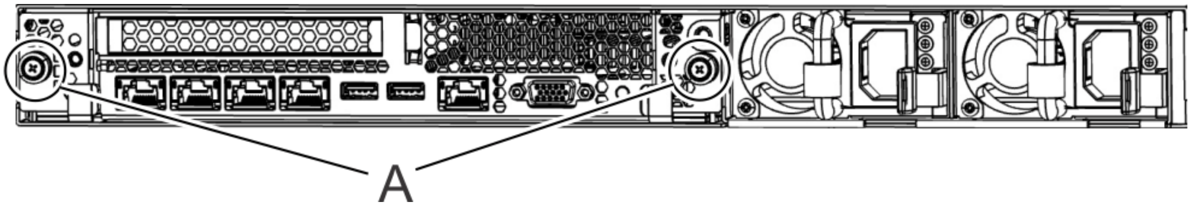

Figure 20. Removing the system backplane screws

d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

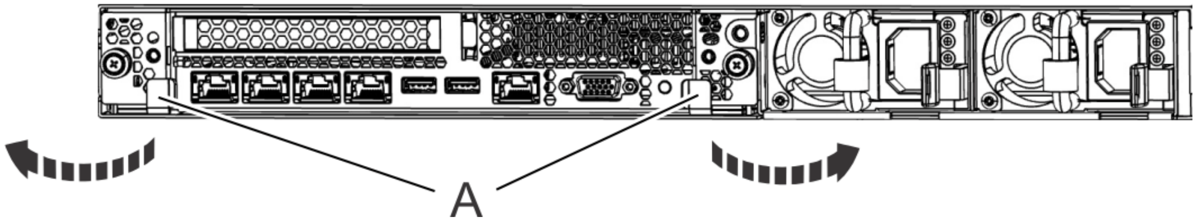

Figure 21. Unlatching the system backplane

e) Support the system backplane by the bottom as you slide it from the system.

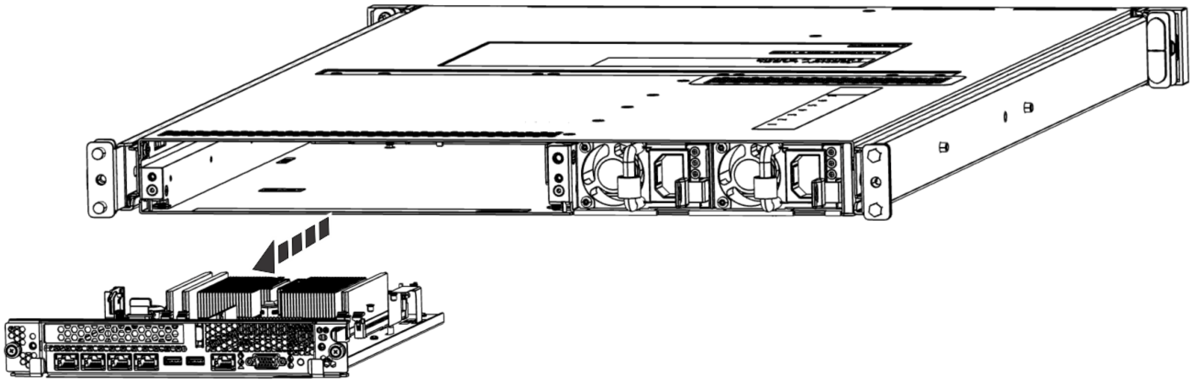

Figure 22. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

4. Remove the memory DIMM.

a) Locate the memory DIMM that you want to remove.

b) Unlock the memory DIMM by simultaneously pushing the locking tabs away from the memory DIMM. Be sure to unlock both tabs at the same time. The lever action of opening the tabs pushes the memory DIMM out of the slot.

c) Hold the memory DIMM by the edges and pull it out of the slot.

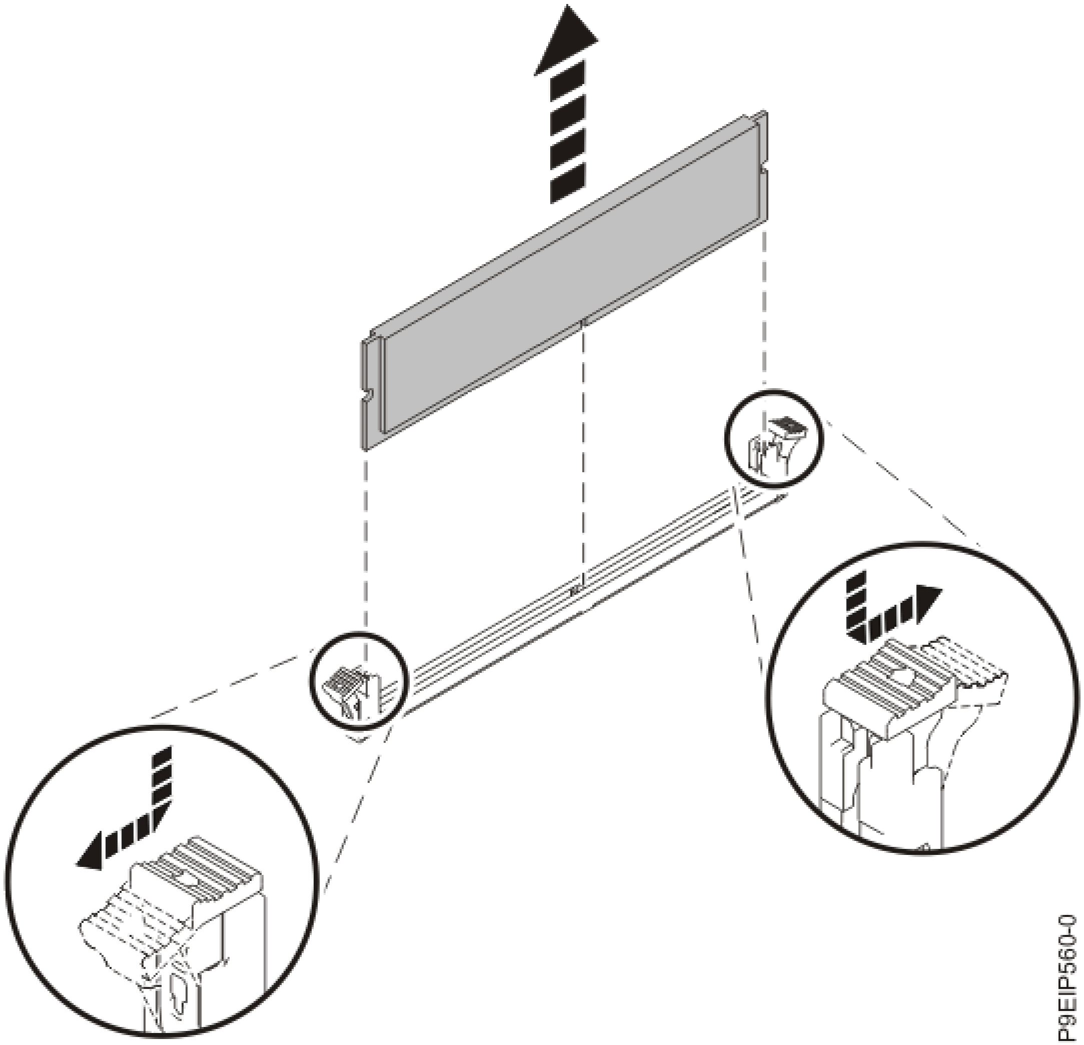

Figure 23. Removing the memory DIMM

5. Insert the memory DIMM.

   a) Grasp the memory DIMM along its edges and align it with the slot on the system backplane.

   ⚠ **Attention:** Memory is keyed to prevent it from being installed incorrectly. Note the location of the key tab within the memory connector before you attempt to install it.

   b) Press firmly on each side of the memory DIMM until the locking tab locks in place with an audible click.

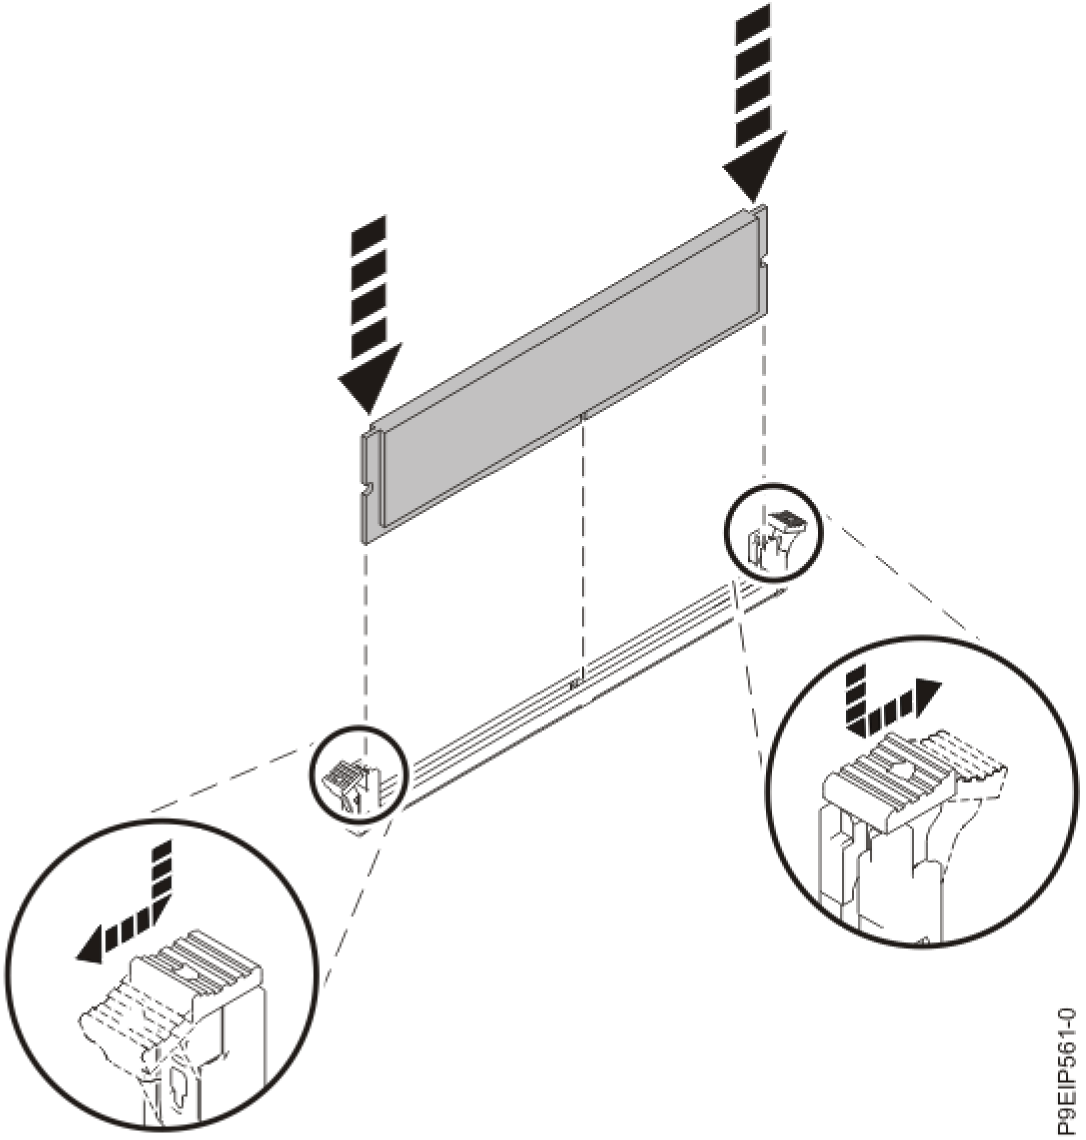

Figure 24. Inserting the memory DIMM

6. Replace the system backplane into the rear of the system.
   a) Ensure that the two system backplane levers are open.
   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

      Important:

      - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
      - Ensure that the system backplane is fully seated and is all the way into the system.
      - You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

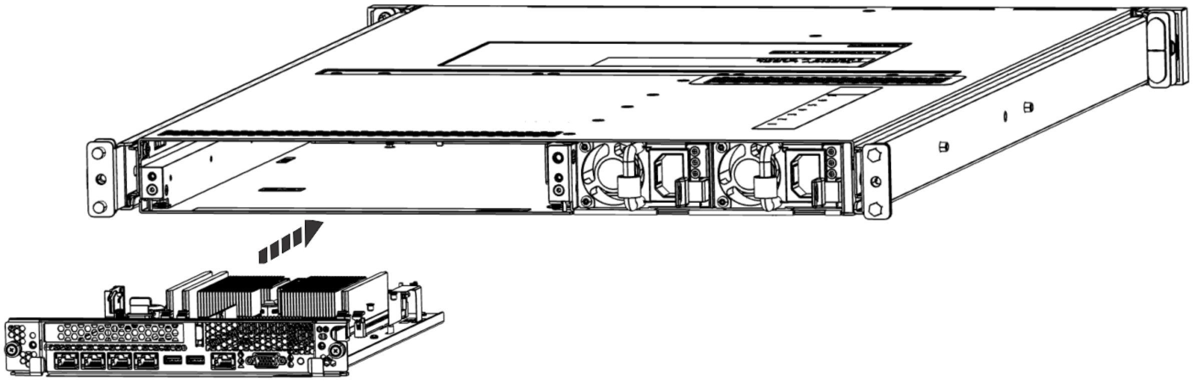

Figure 25. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

7. Power on the system for operation.

   For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the control panel in the 7063-CR2

Learn how to remove and replace the control panel in the IBM Power Systems HMC (7063-CR2) system.

Removing the control panel from the 7063-CR2 system

To remove the control panel from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

2. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.

3. Place the system in the service position on an ESD surface on a table.

   For instructions, see "Placing a 7063-CR2 system into the service position" on page 101.

4. Turn the system upside down on the ESD surface.
5. Remove the two screws **(A)** that secure the control panel to the bottom of the system.

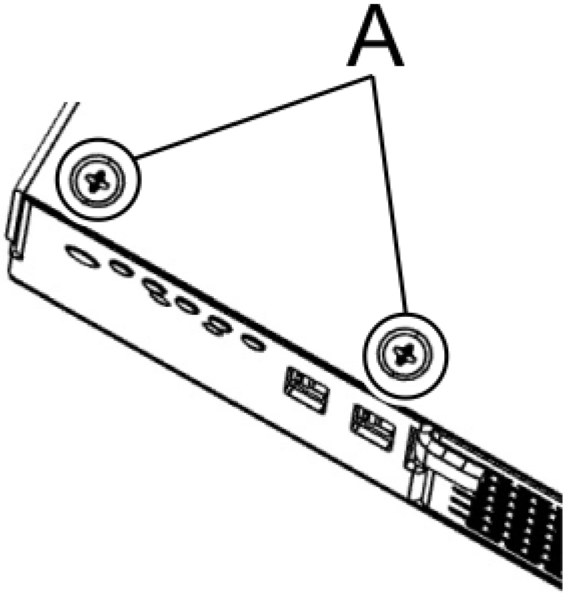

Figure 26. Removing the control panel bottom screws

6. Turn the system right side up on the ESD surface.
7. Remove the screw **(A)** that secures the control panel to the right side of the system.

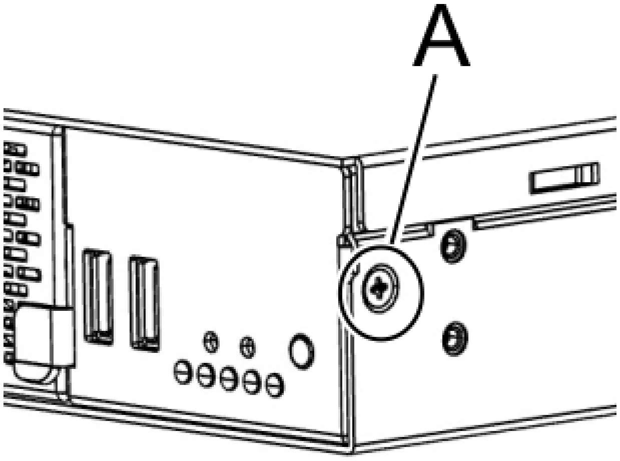

Figure 27. Removing the control panel side screw

8. Remove the fan cover from the front of the system.
   a) Rotate the two levers **(A)** on each side of the fan cover up and out to unlock the fan cover.
   b) Pull the cover away from the system.

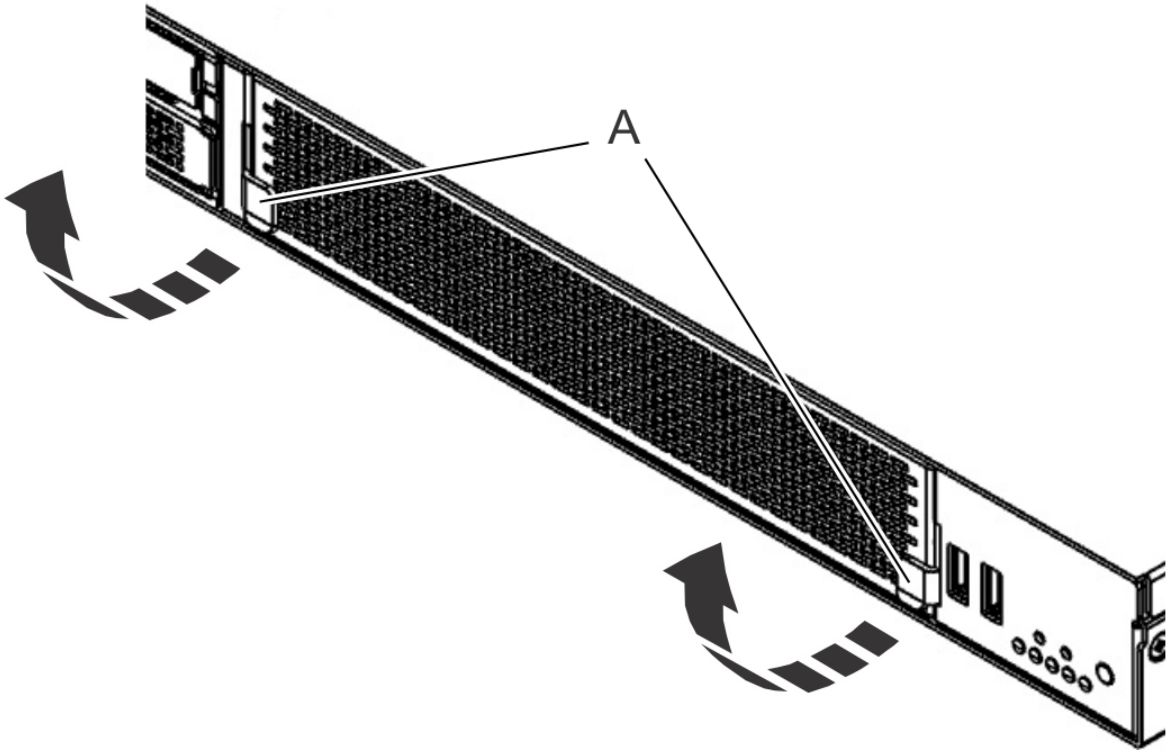

Figure 28. Removing the fan cover

9. Slide the control panel out of the system. Be careful with the cables that are attached to the control panel.

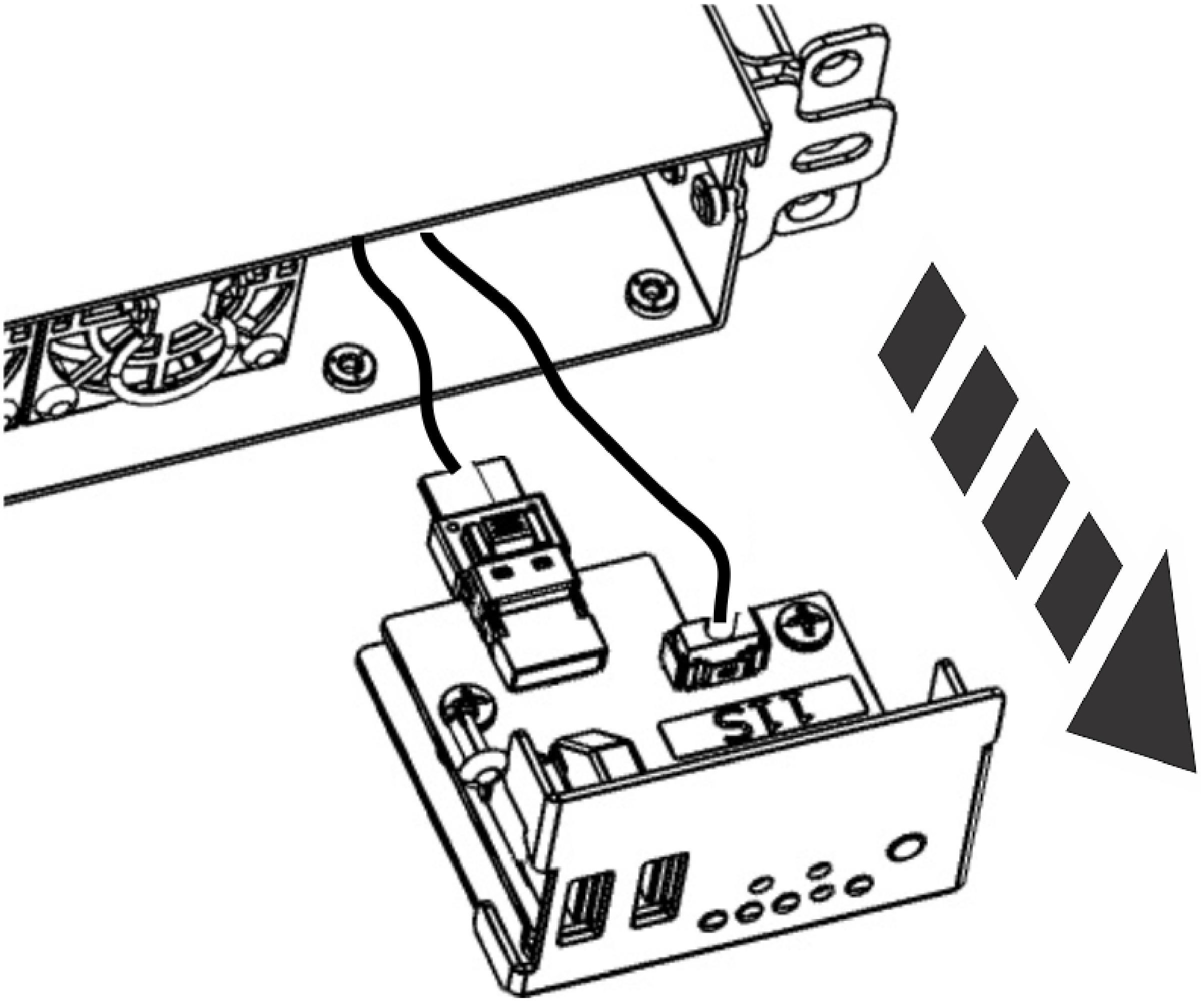

Figure 29. Removing the control panel from the system

10. Disconnect the USB cable and the control panel cable from the control panel.

    Use your thumb or finger to press the release latch on the connector to remove the cable.

11. Remove the three screws that secure the control panel card to the control panel cover.

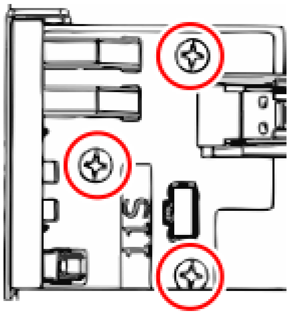

Figure 30. Removing the control panel screws

Replacing the control panel in the 7063-CR2 system

To replace the control panel in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Remove the new control panel card from new control panel cover.
3. Align the new control panel card on the current control panel cover.

   The current control panel cover has the system serial number and needs to be used.
4. Replace the three screws that secure the control panel card to the control panel cover.

   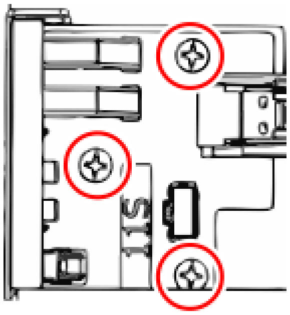

   Figure 31. Replacing the control panel screws
5. Connect the USB cable and the control panel cable to the control panel.
6. Install the control panel into the system.

   Ensure that the USB and power switch cables are not pinched when your insert the control panel.

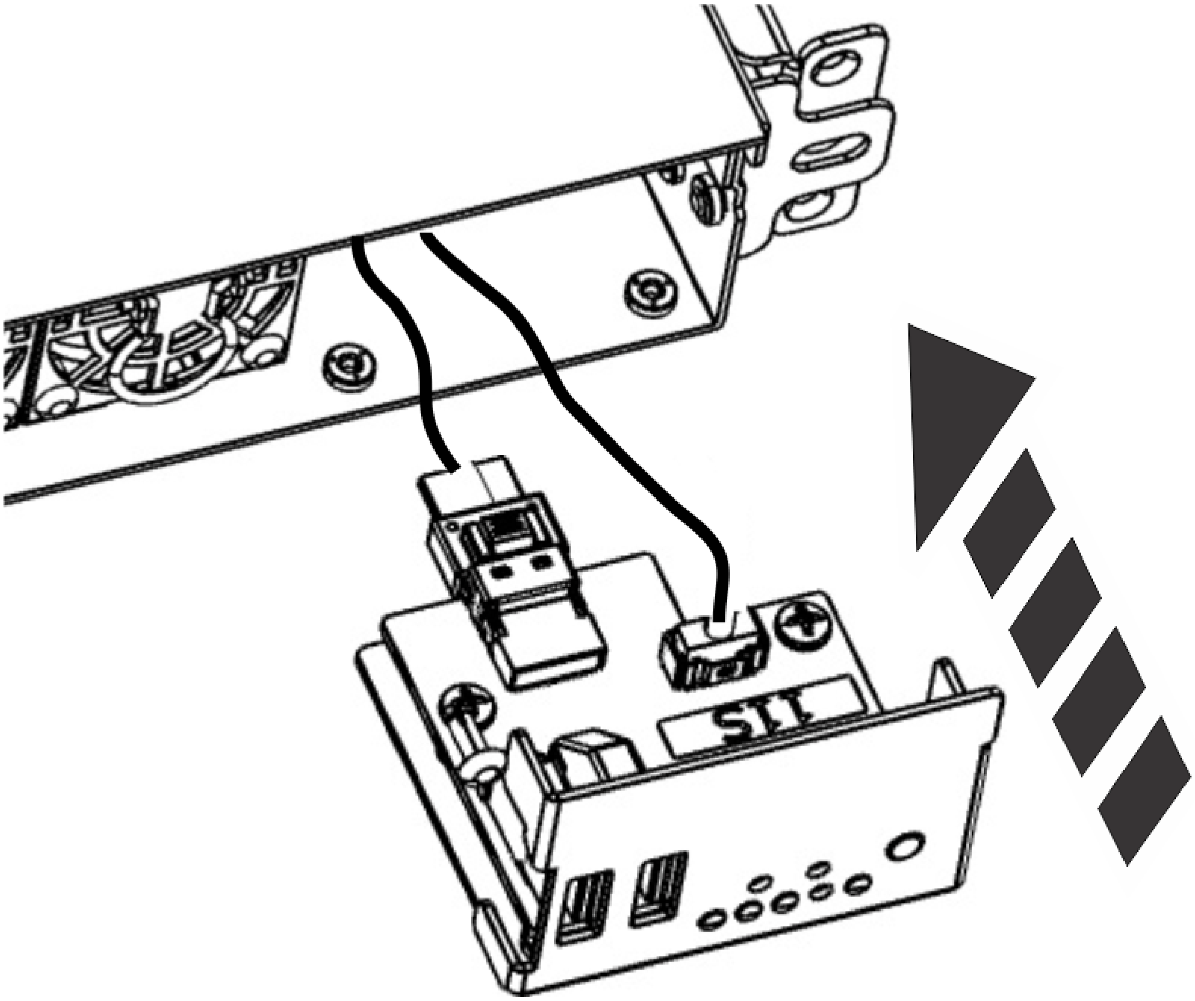

Figure 32. Installing the control panel into the system

7. Install the screw **(A)** to secure the control panel to the right side of the system.

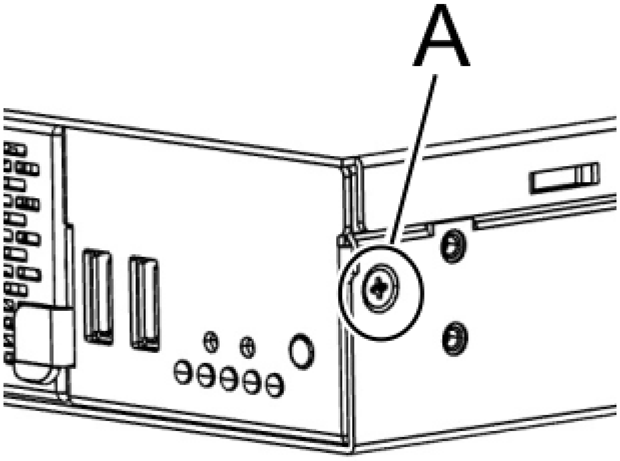

Figure 33. Installing the control panel side screw

8. Turn the system upside down on the ESD surface.
9. Install the two screws **(A)** to secure the control panel to the bottom of the system.

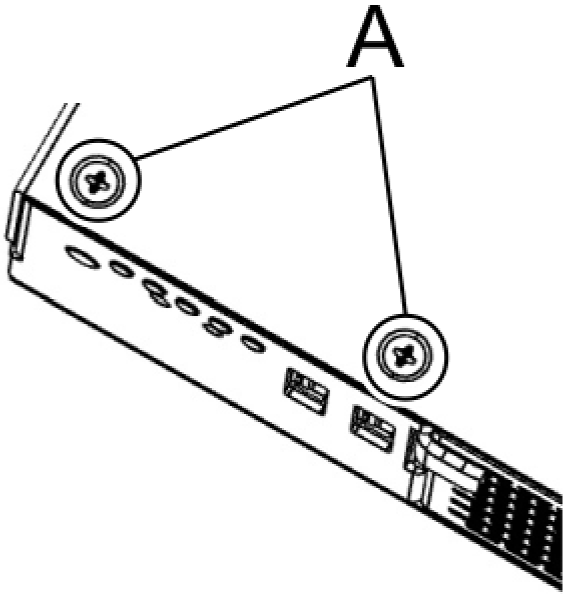

Figure 34. Installing the control panel bottom screws

10. Turn the system right side up on the ESD surface.
11. Replace the fan cover.
    a) Ensure that the two levers are open.
    b) Put the fan cover in place.
    c) Rotate the two levers on each side of the fan cover down and in to secure the cover to the system.

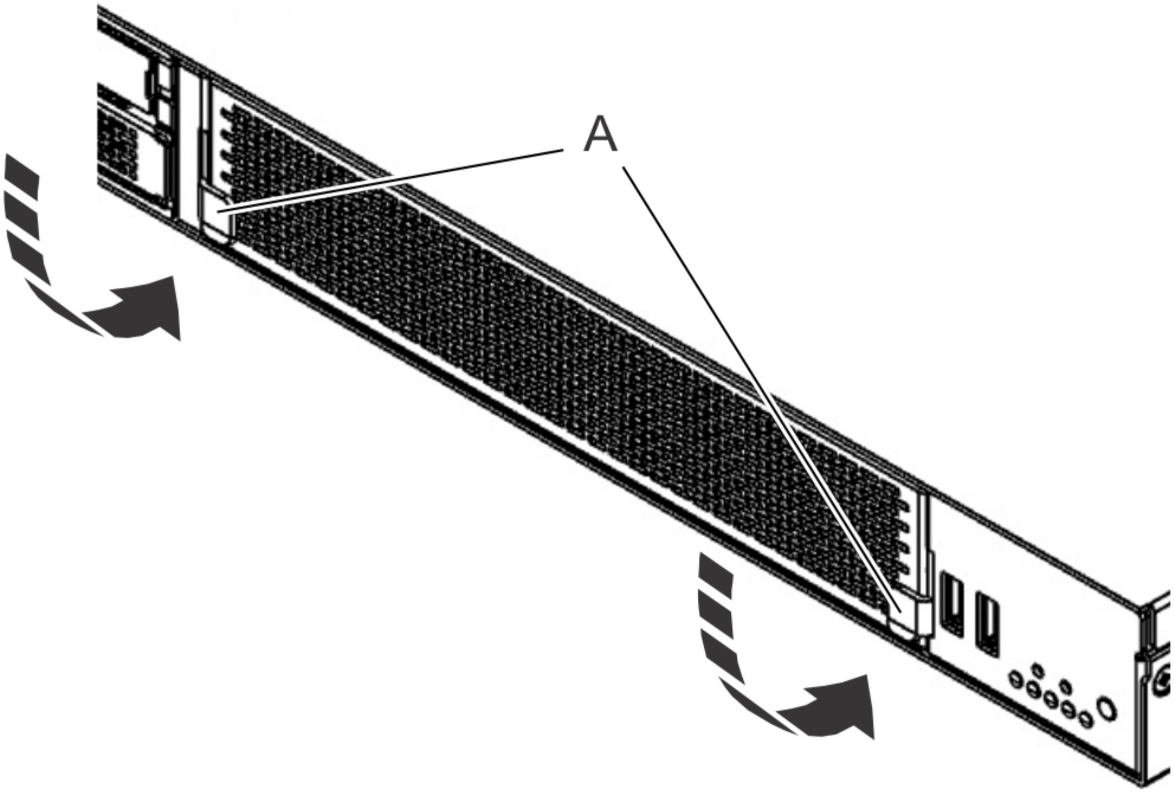

Figure 35. Replacing the fan cover

12. Replace the system in the rack and replace the components that you removed.
    For instructions, see "Placing a 7063-CR2 system into the operating position" on page 103.

13. Power on the system for operation.

    For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing PCIe adapters in the 7063-CR2

Learn how to remove and replace Peripheral Component Interconnect (PCI) Express (PCIe) adapters in the IBM Power Systems HMC (7063-CR2) system.

About this task

Attention: For safety and airflow purposes, if you remove a PCIe adapter from the system, you must ensure that a PCIe filler and a PCIe tail-stock filler are present.

Parts must be replaced with the identical part in the exact same place.

Removing a PCIe adapter from the 7063-CR2 system

To remove a PCIe adapter from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

Attention: For safety and airflow purposes, if you remove a PCIe adapter from the system, you must ensure that a PCIe filler and a PCIe tail-stock filler are present.

Parts must be replaced with the identical part in the exact same place.

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.
2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   Attention:

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.
3. Remove the system backplane from the rear of the system.

   a) Label and remove the two power cables.

      For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.

   b) Label and remove the signal cables from the rear of the system.

   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

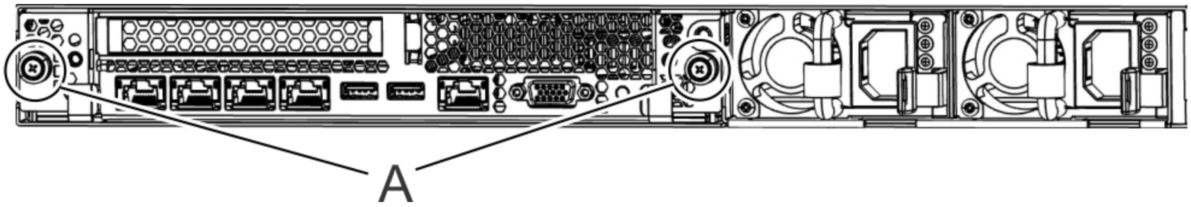

Figure 36. Removing the system backplane screws

d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

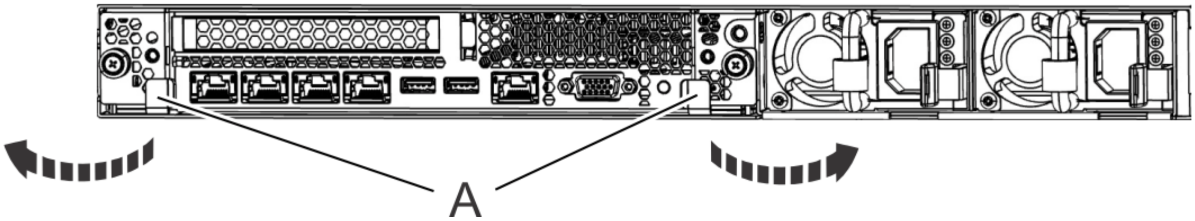

Figure 37. Unlatching the system backplane

e) Support the system backplane by the bottom as you slide it from the system.

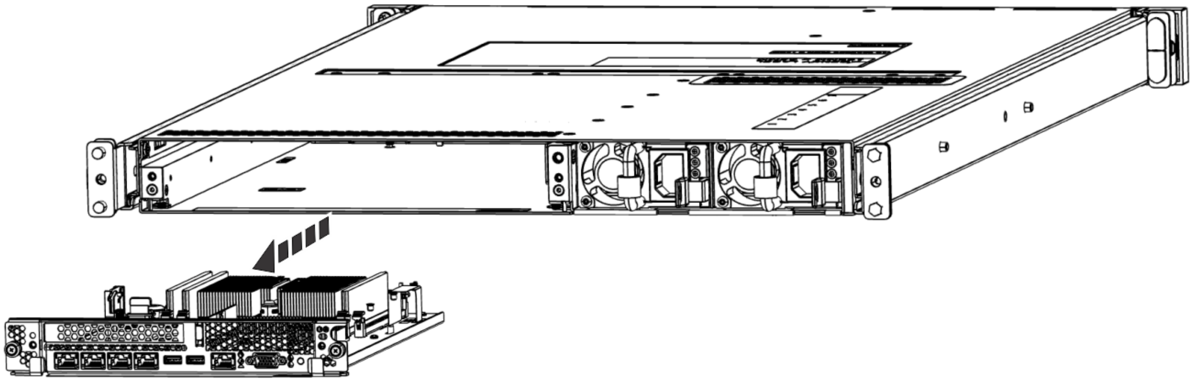

Figure 38. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

4. Remove the PCIe adapter from the PCIe riser.

a) Remove the screw that secures the PCIe adapter to the system backplane.

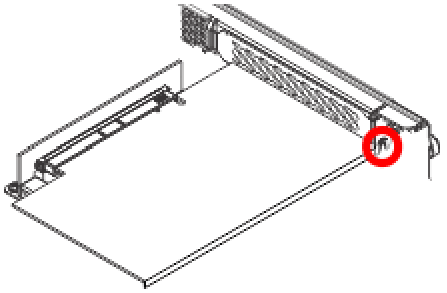

Figure 39. Removing the tail-stock screw

b) Open the retainer clip that secures the PCIe adapter to the PCIe riser, by moving the blue lever to the unlocked position.

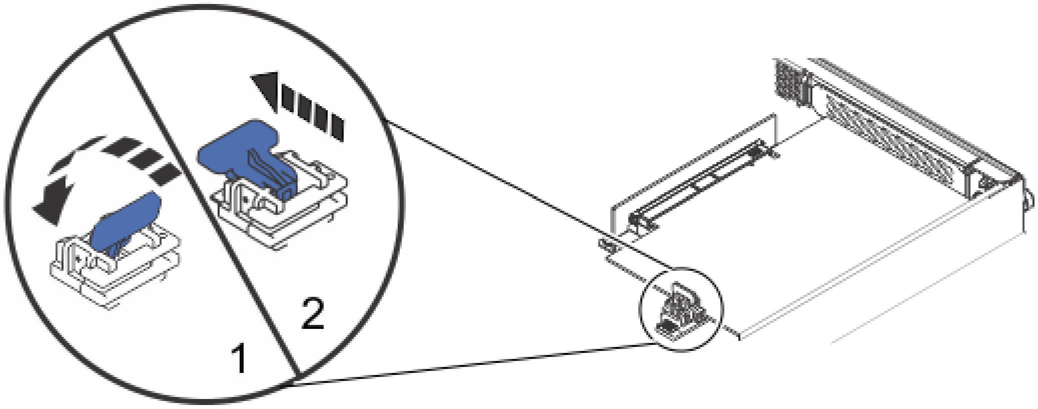

Figure 40. Removing the PCIe adapter retaining latch

c) Move the retainer clip away from the PCIe adapter.

d) Remove the PCIe adapter from the PCIe riser.

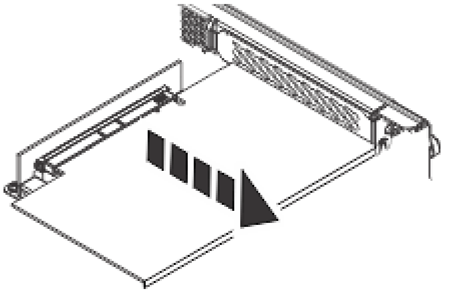

Figure 41. Removing the PCIe adapter

Replacing a PCIe adapter in the 7063-CR2 system

To replace a PCIe adapter in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Replace the PCIe adapter.

a) Replace the PCIe adapter into the PCIe riser.

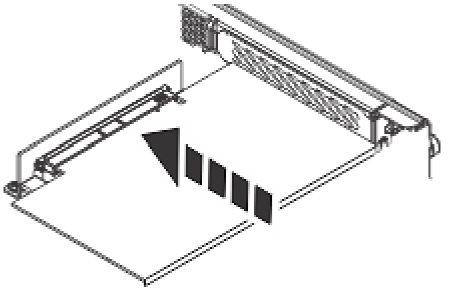

Figure 42. Replacing the PCIe adapter

b) Replace the screw that secures the PCIe adapter to the system backplane.

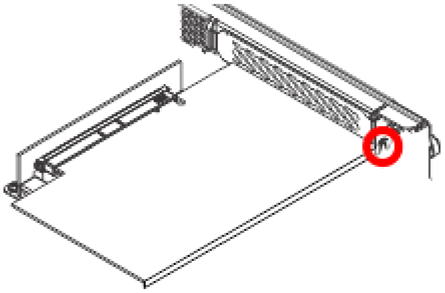

Figure 43. Replacing the tail-stock screw

c) Move the retainer clip to secure the PCIe adapter to the PCIe riser. Ensure that the clip is fully seated around the edge of the adapter.

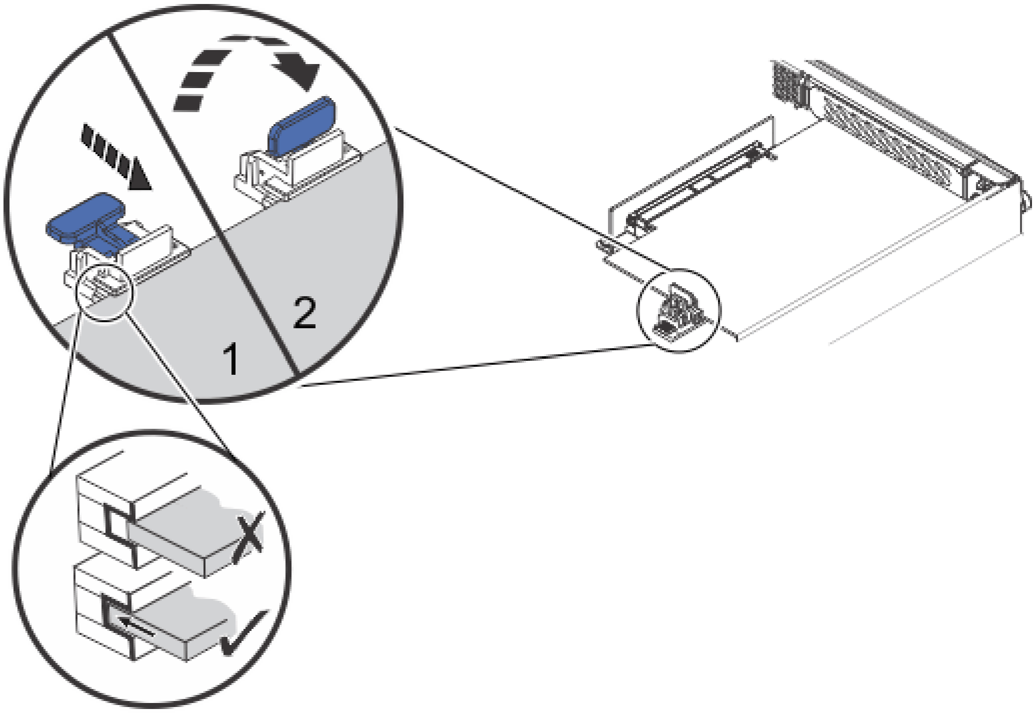

Figure 44. Replacing the PCIe adapter retaining latch

d) Close the retainer clip that secures the PCIe adapter to the PCIe riser.

3. Replace the system backplane into the rear of the system.

   a) Ensure that the two system backplane levers are open.

   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

   Important:

   - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
   - Ensure that the system backplane is fully seated and is all the way into the system.
   - You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

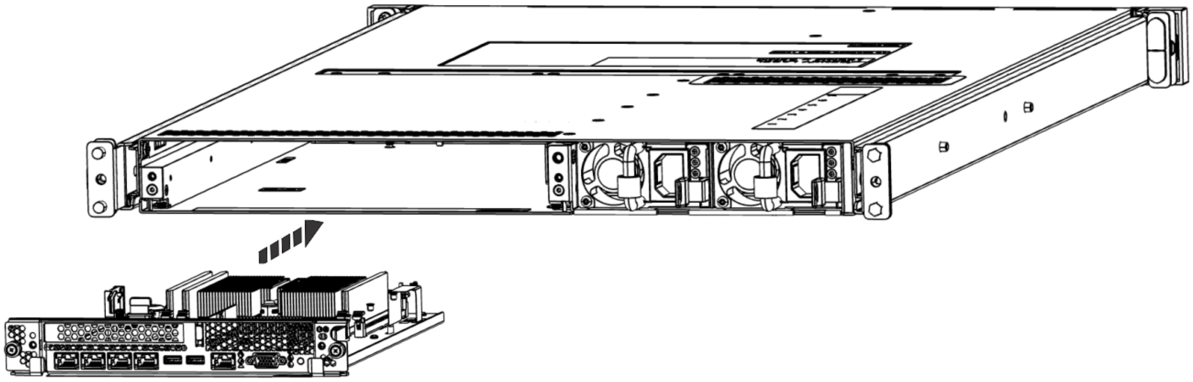

Figure 45. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

4. Power on the system for operation.

   For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the PCIe riser in the 7063-CR2

Learn how to remove and replace the PCIe riser in the IBM Power Systems HMC (7063-CR2) system.

Removing the PCIe riser from the 7063-CR2 system

To remove the PCIe riser from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.

2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   Attention:

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

3. Remove the system backplane from the rear of the system.

   a) Label and remove the two power cables.

For instructions, see “Disconnecting the power cords from a 7063-CR2 system” on page 100.

b) Label and remove the signal cables from the rear of the system.

c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

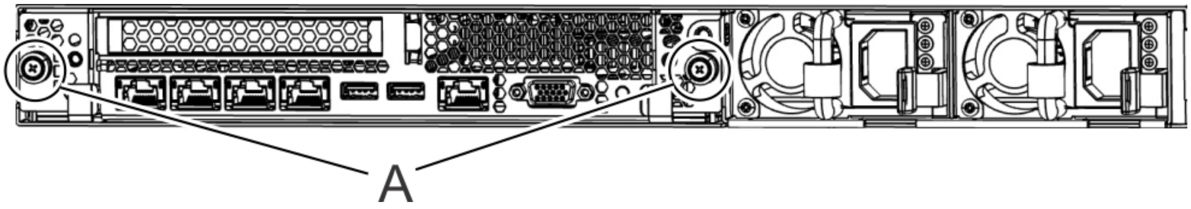

Figure 46. Removing the system backplane screws

d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

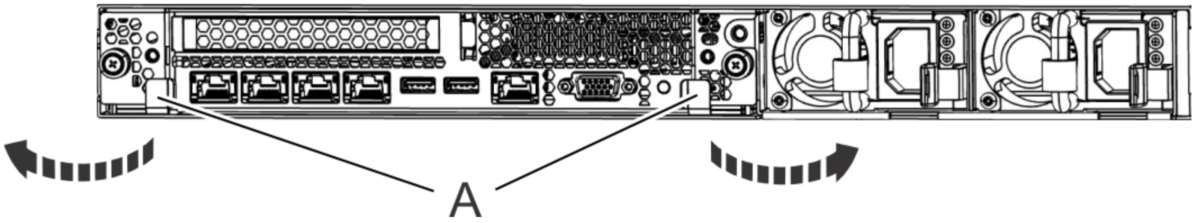

Figure 47. Unlatching the system backplane

e) Support the system backplane by the bottom as you slide it from the system.

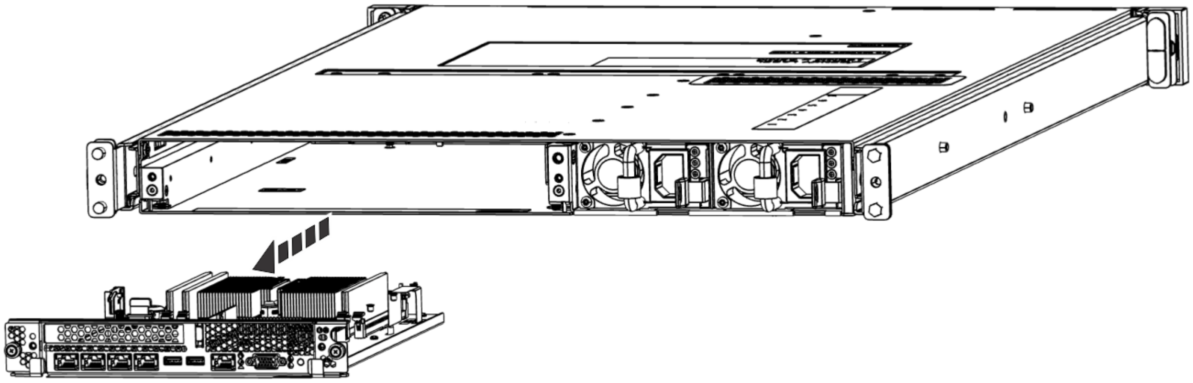

Figure 48. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

4. Remove the PCIe adapter or filler from the PCIe riser.

   For instructions, see “Removing a PCIe adapter from the 7063-CR2 system” on page 33.

5. Remove the four screws that secure the PCIe riser to the system backplane.

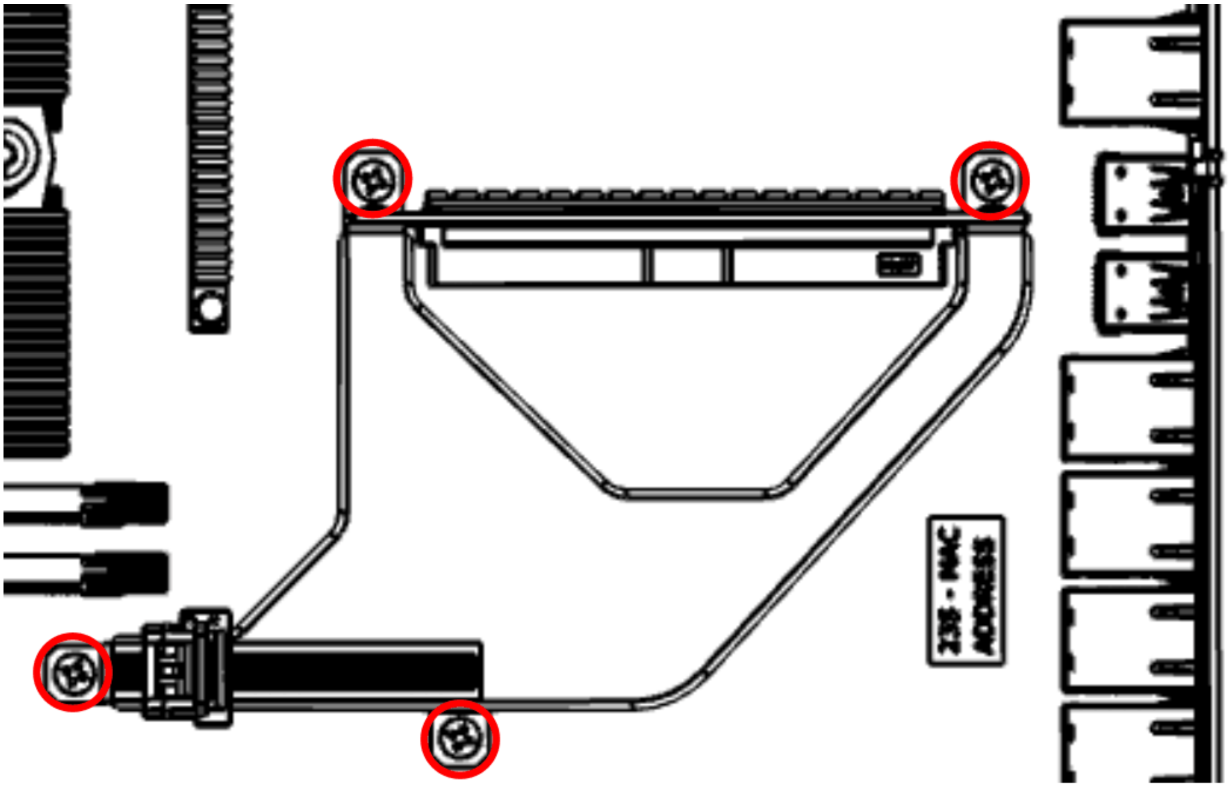

Figure 49. Removing the PCIe riser screws

6. Lift the PCIe riser from the system and place it on an ESD surface.

Replacing the PCIe riser in the 7063-CR2 system

To replace the PCIe riser in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Install the PCIe riser into the system. Insert the PCIe riser into the system backplane until the riser card is fully seated in the socket on the backplane.
3. Replace the four screws to secure the PCIe riser to the system backplane.

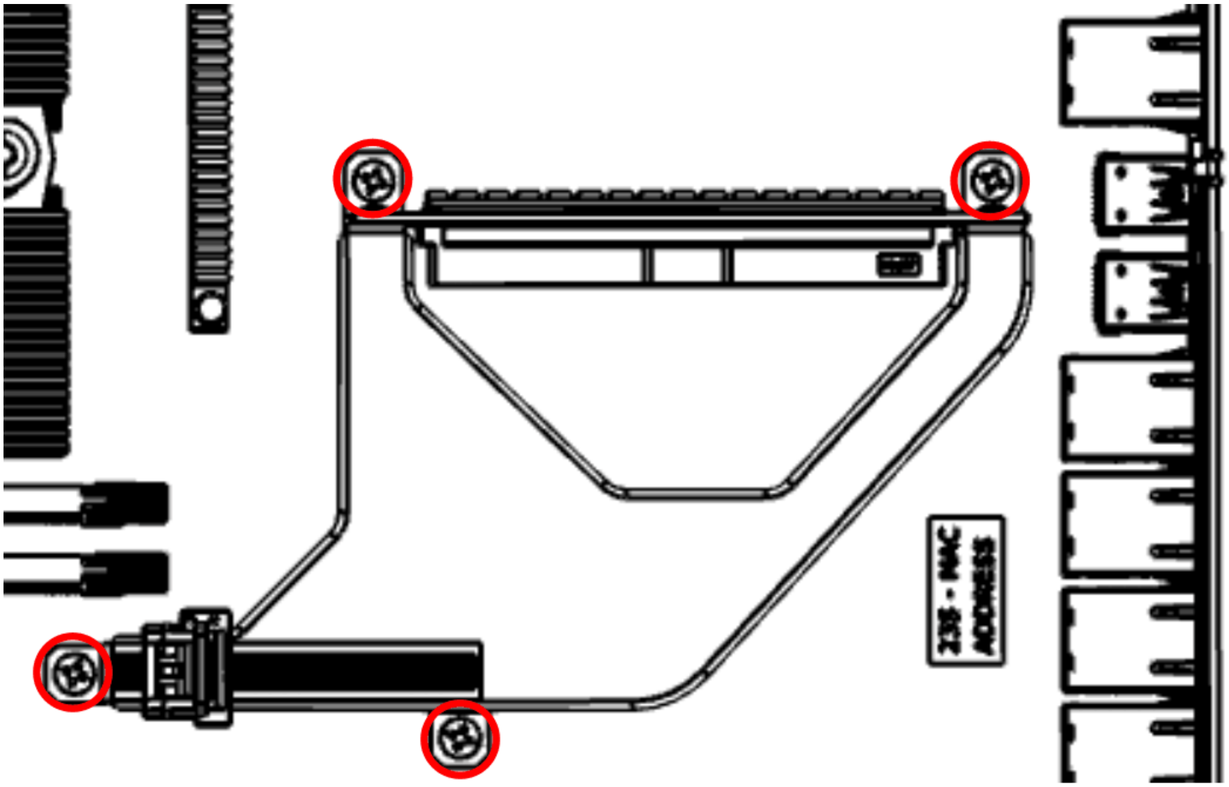

Figure 50. Replacing the PCIe riser screws

4. Replace the PCIe adapter or filler.

   For instructions, see "Replacing a PCIe adapter in the 7063-CR2 system" on page 35.

5. Replace the system backplane into the rear of the system.

   a) Ensure that the two system backplane levers are open.

   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

      Important:

      - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
      - Ensure that the system backplane is fully seated and is all the way into the system.
      - You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

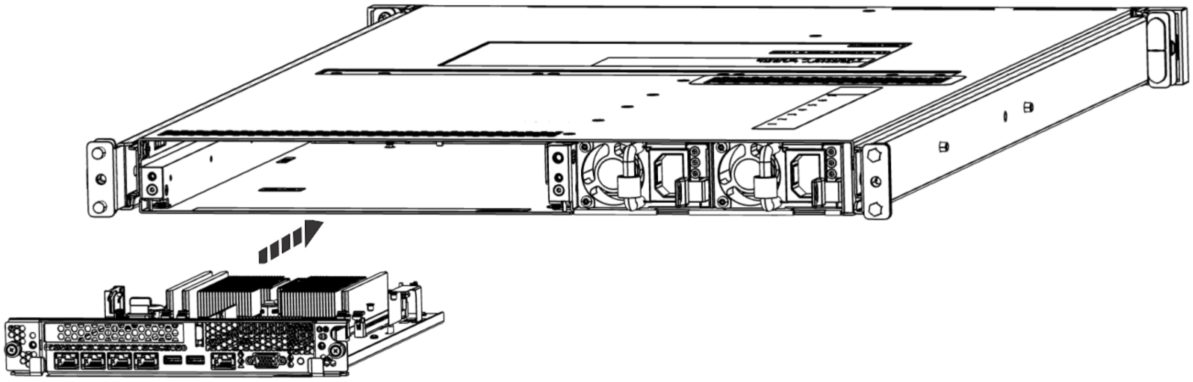

Figure 51. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

6. Power on the system for operation.

For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the power distribution board in the 7063-CR2

Learn how to remove and replace the power distribution board in the IBM Power Systems HMC (7063-CR2) system.

Removing the power distribution board from the 7063-CR2 system

To remove the power distribution board from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

You can use a magnetic tip screwdriver to remove and replace the screws.

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.

2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

3. Remove all the fans from the system.

   For instructions, see "Removing a fan from the 7063-CR2 system" on page 18.

4. Remove the system backplane from the rear of the system.

   a) Label and remove the two power cables.

   For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.

   b) Label and remove the signal cables from the rear of the system.

   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

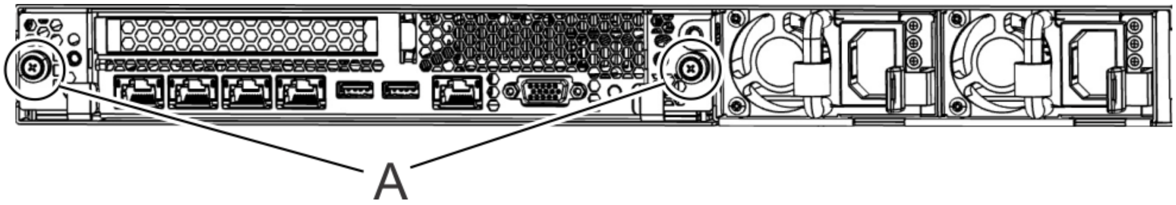

Figure 52. Removing the system backplane screws

d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

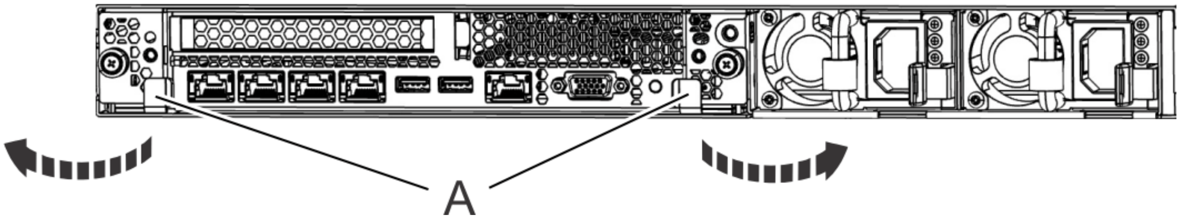

Figure 53. Unlatching the system backplane

e) Support the system backplane by the bottom as you slide it from the system.

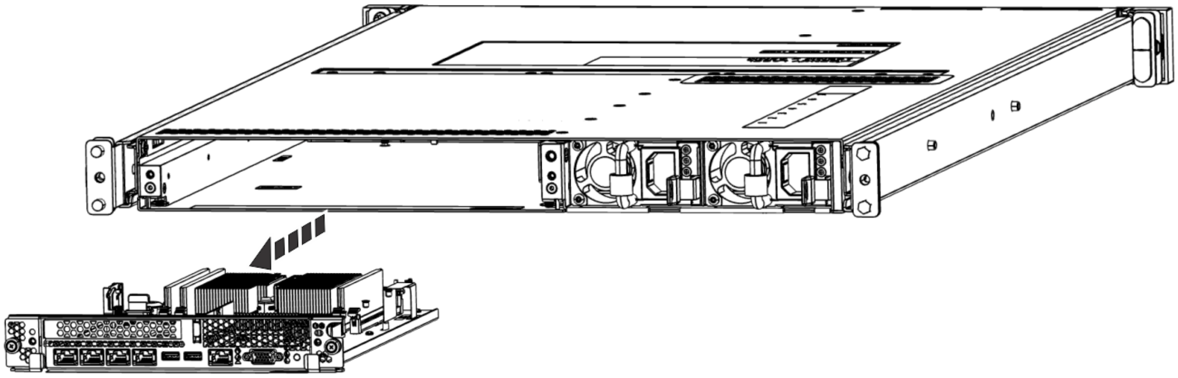

Figure 54. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

5. Place the system in the service position on an ESD surface on a table.

   For instructions, see “Placing a 7063-CR2 system into the service position” on page 101.

6. Remove the service access cover.

   For instructions, see “Removing the service access cover from a 7063-CR2 system” on page 104.

7. Label and unplug the five cables **(A)** from the power distribution board. Unplug the drive cables from the rear of the drive holder.

   Use your thumb or finger to press the release latch on the connector to remove the cable.

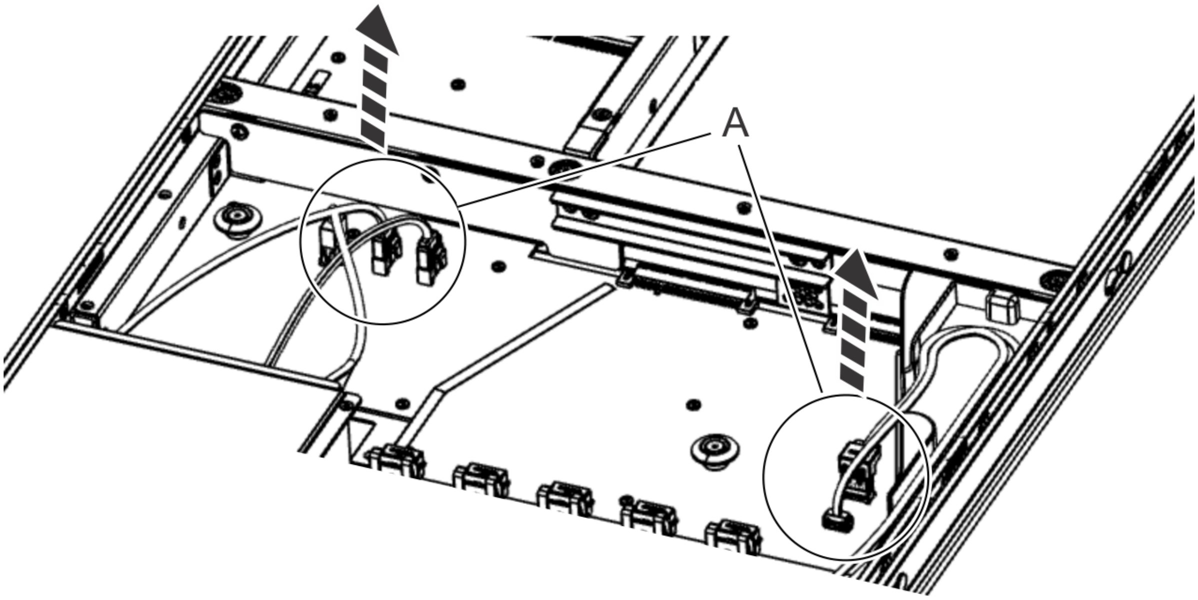

Figure 55. Removing cables from the power distribution board

8. Remove the right plastic air baffle **(A)** from the right edge of the system; lift it straight up.

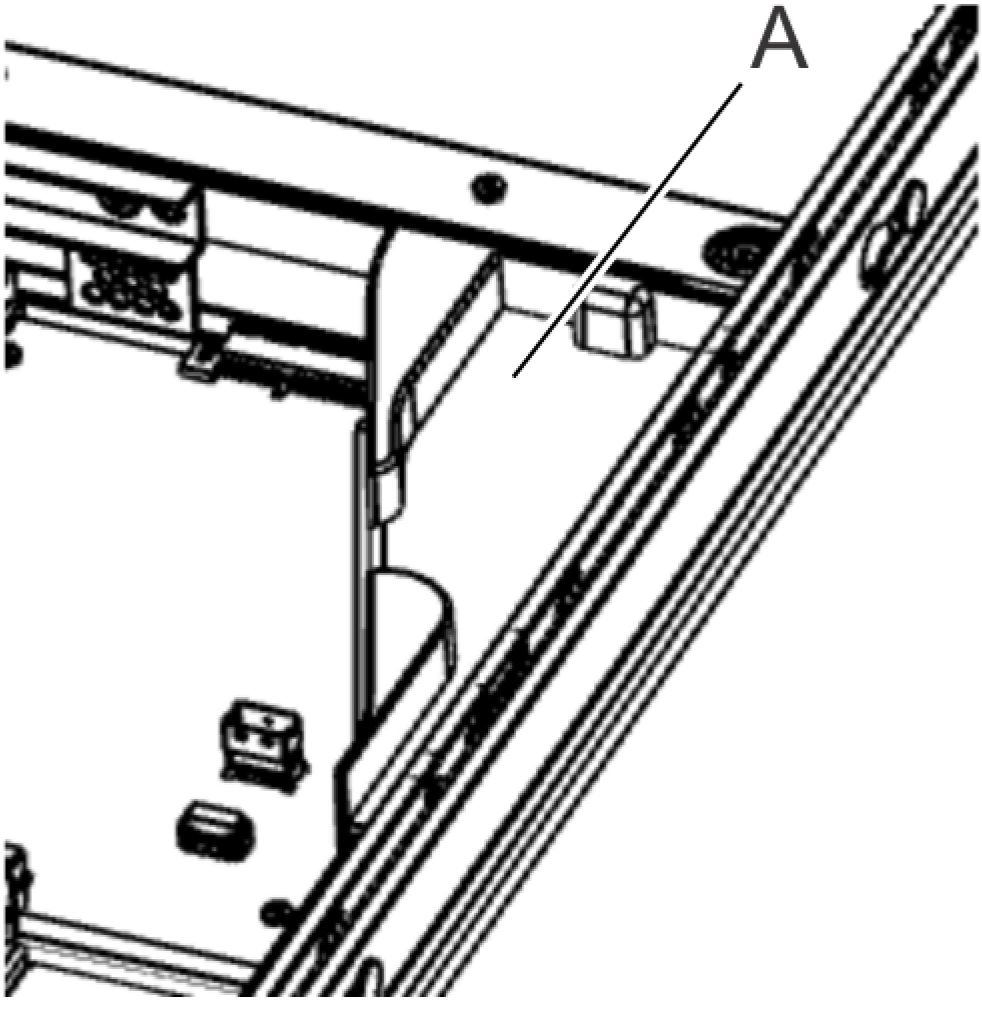

Figure 56. Right air baffle

9. Remove the cross bar from inside the system.

a) Remove the five screws from the front of the cross bar.

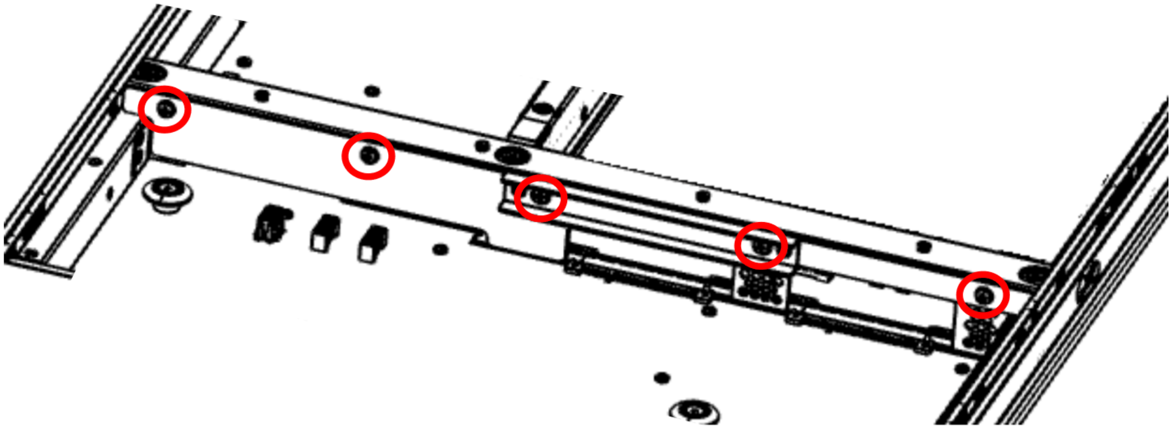

Figure 57. Cross bar front screws

b) Remove the air baffle **(A)** from the cross bar.

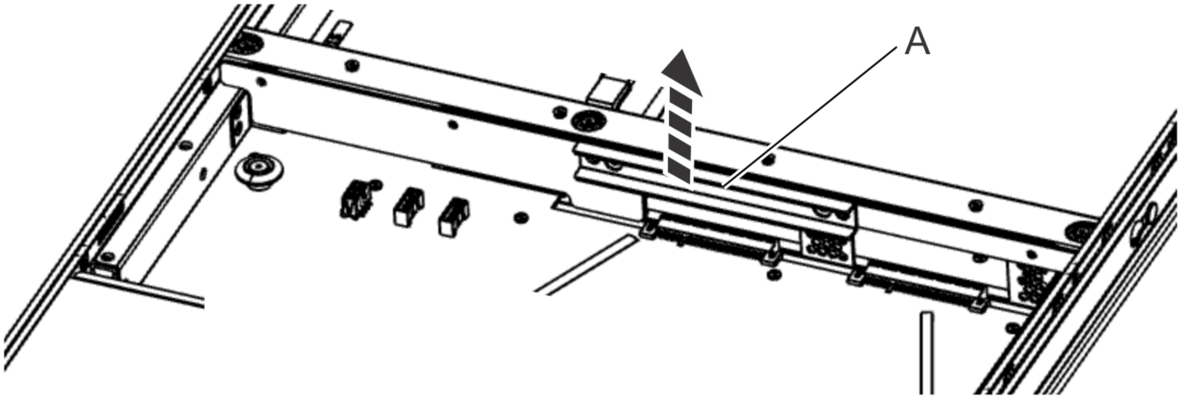

Figure 58. Removing the air baffle

c) Loosen the three screws at the top of the cross bar.

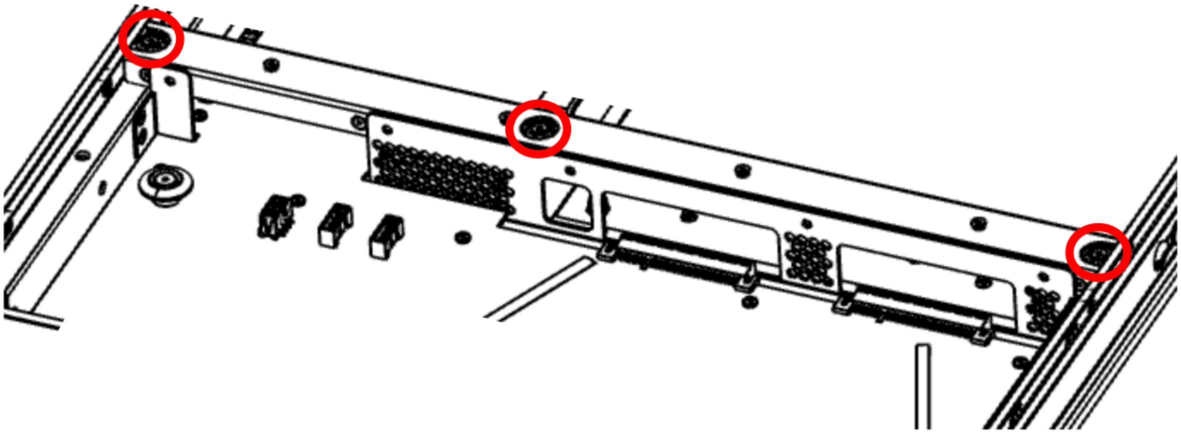

Figure 59. Cross bar top screws

d) Lift the cross bar out of the system.

10. Remove the 19 screws that secure the power distribution board to the system.

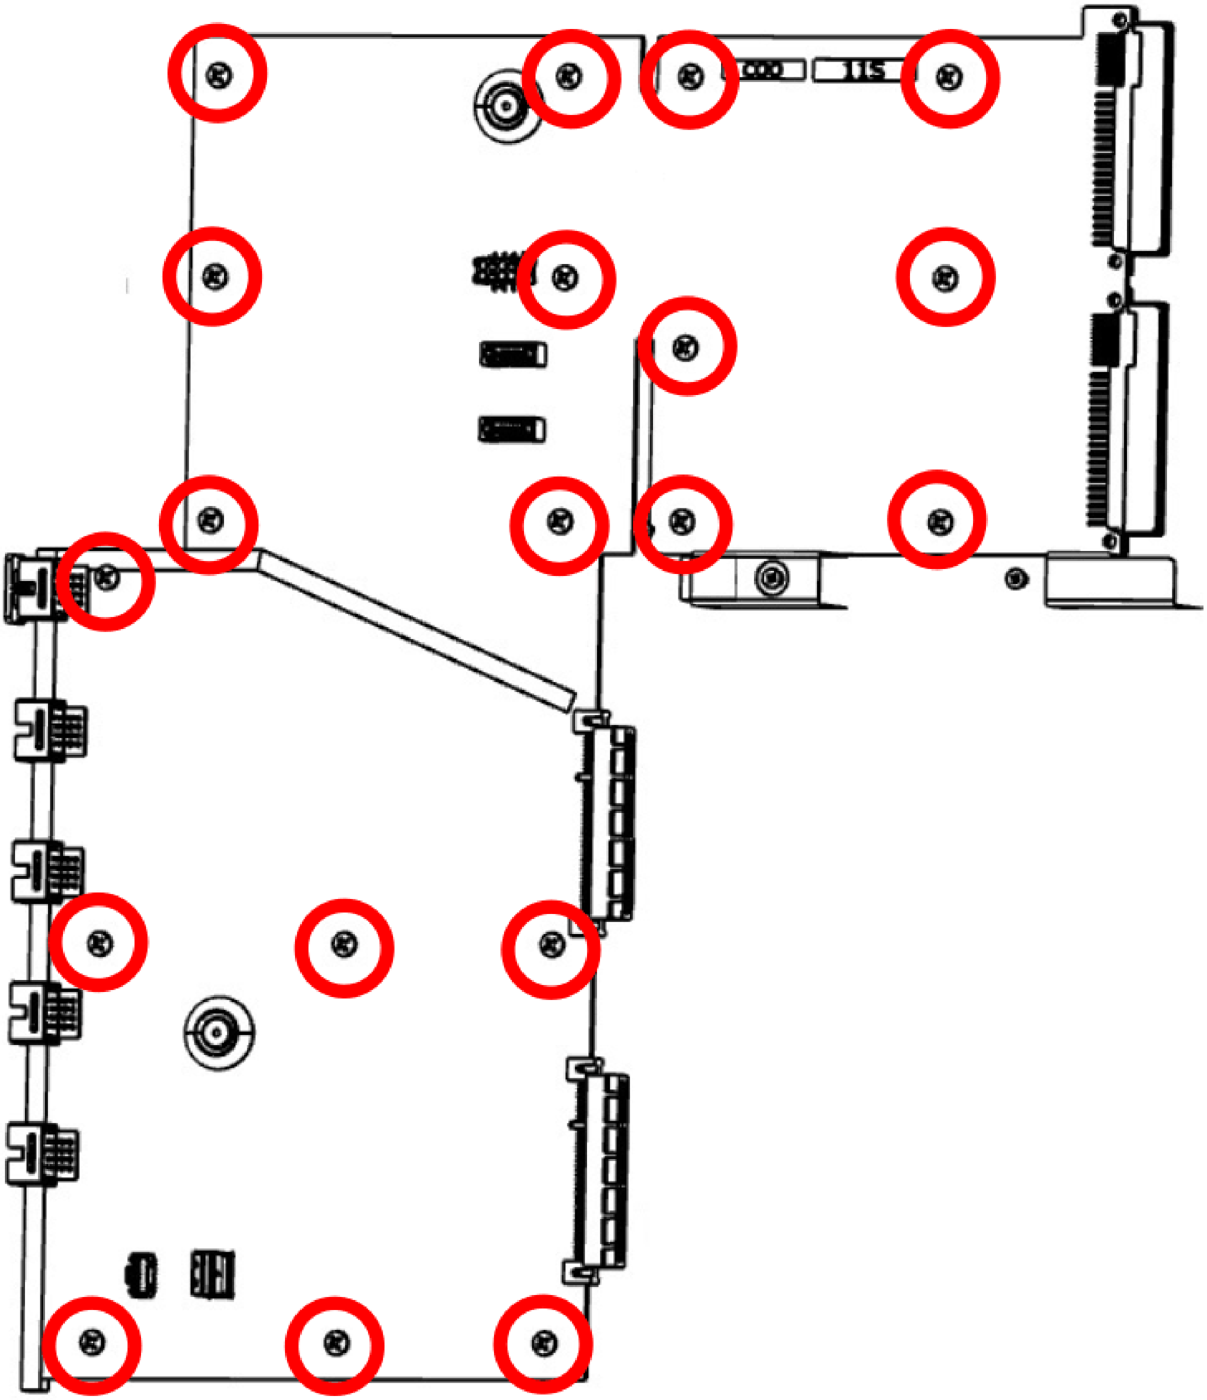

Figure 60. Power distribution board screws

11. Use the two finger grips to lift the power distribution board up and out of the system.
12. Place the power distribution board on an ESD surface.

Replacing the power distribution board in the 7063-CR2 system

To replace the power distribution board in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Use the two finger grips to lift the power distribution board and set the power distribution board into the system.
3. Replace the 19 screws to secure the power distribution board to the system.

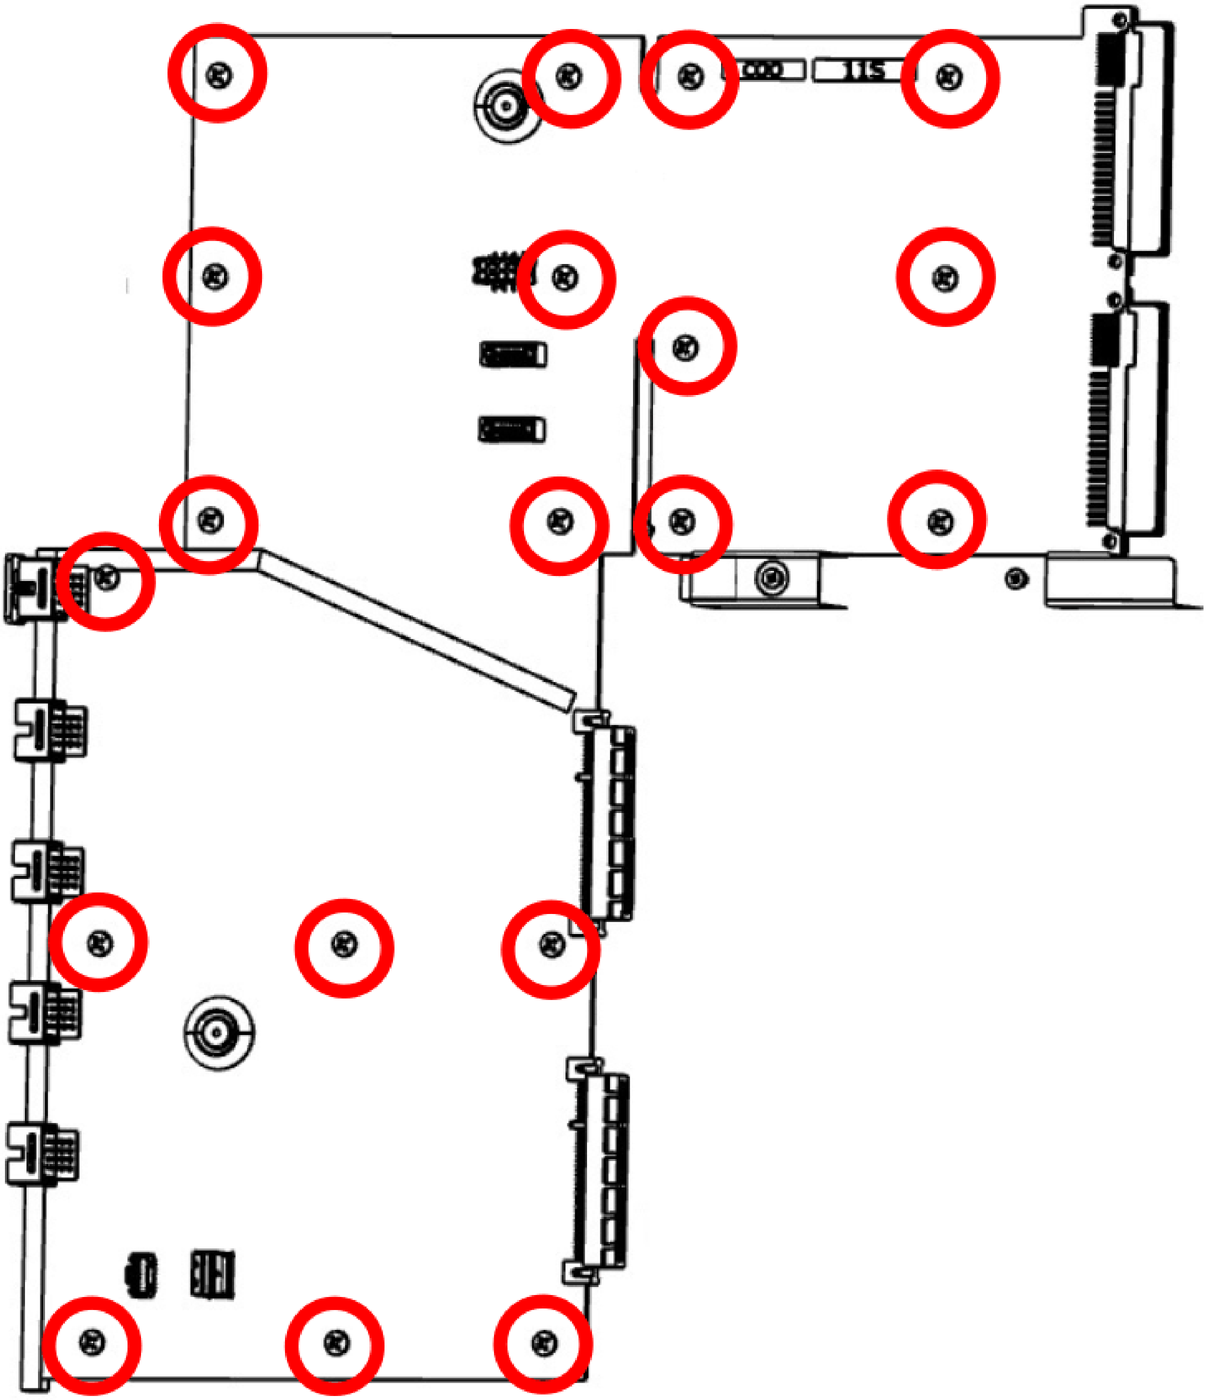

Figure 61. Power distribution board screws

4. Replace the cross bar inside the system.
   a) Put the cross bar into its location inside the system. Ensure that the five screw holes face the front.
   b) Tighten the three screws at the top of the cross bar.

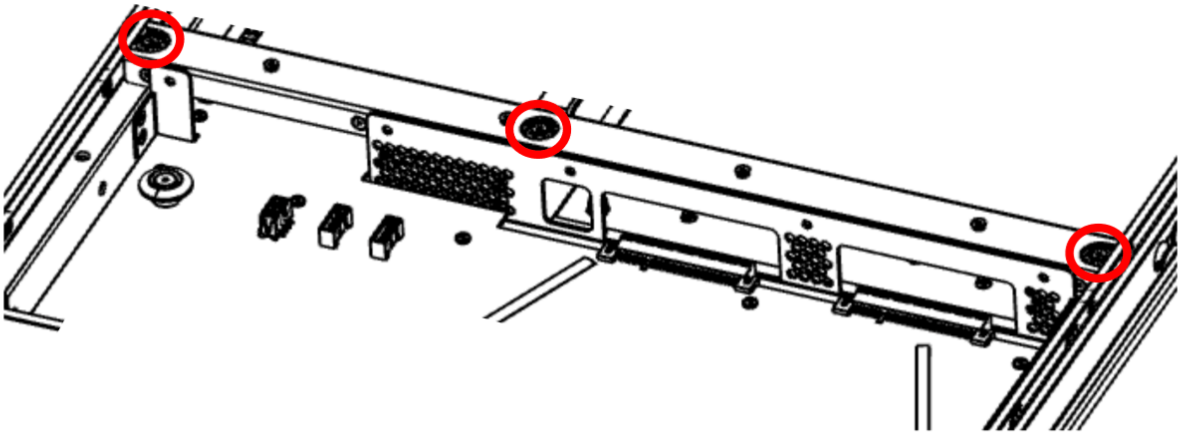

Figure 62. Cross bar top screws

c) Replace the air baffle **(A)** into the system, placing it against the front side of the cross bar.

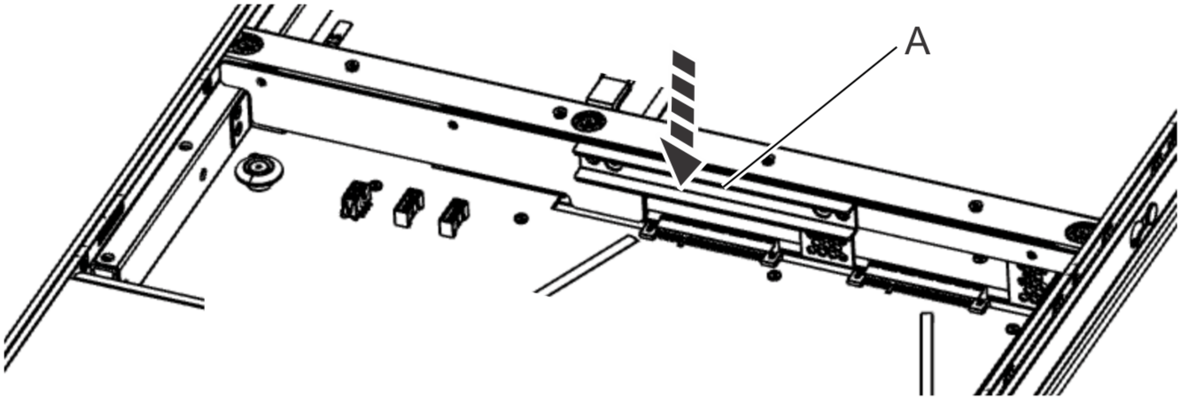

Figure 63. Replacing the air baffle

d) Replace the five screws into the front of the cross bar.

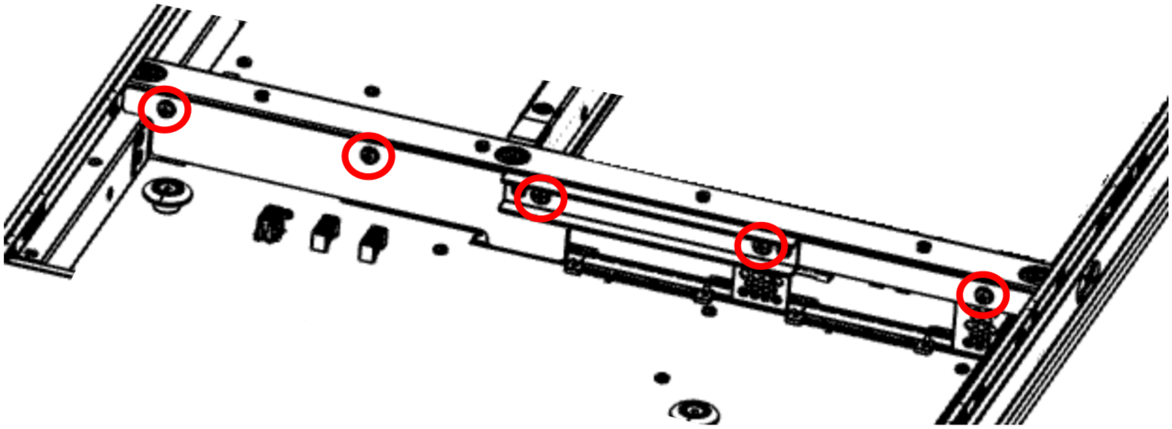

Figure 64. Cross bar front screws

5. Replace the plastic air baffle **(A)** on the right edge of the system; place it straight down.

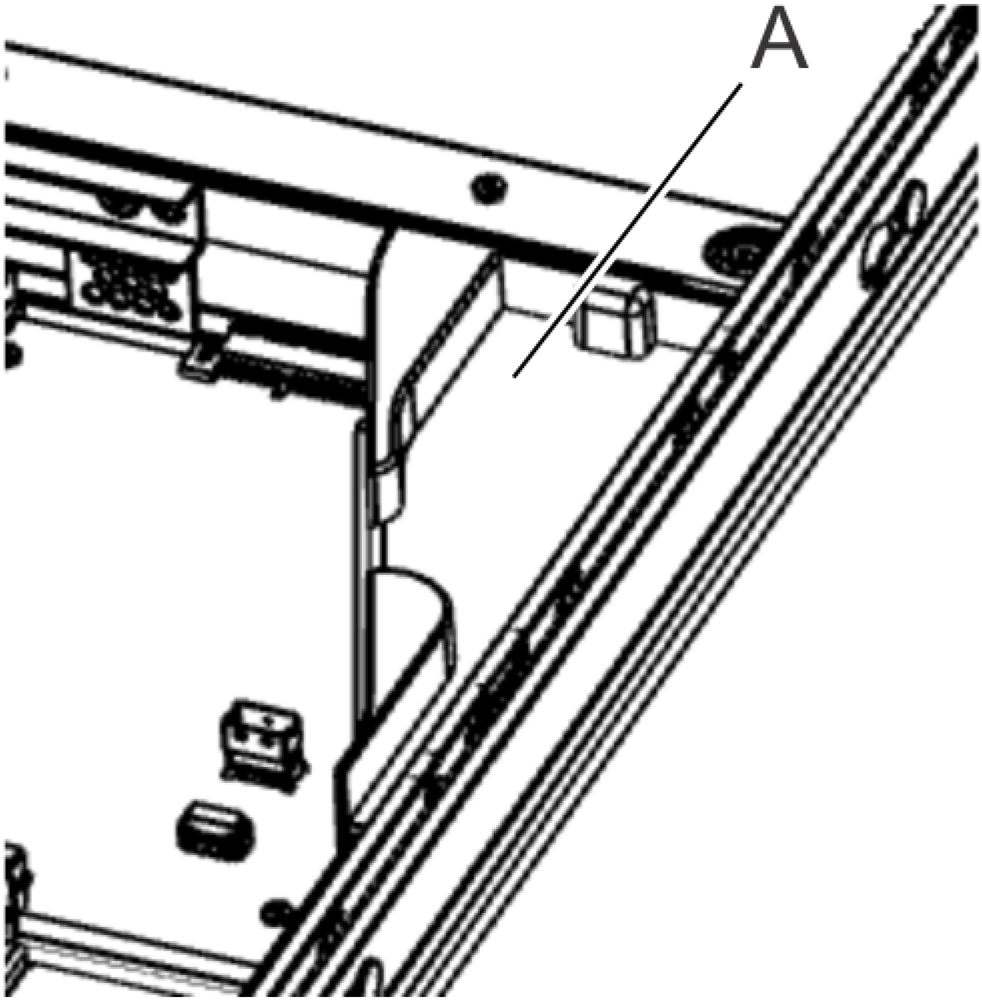

Figure 65. Right air baffle

6. Using your labels, replace the five cables **(A)** into the power distribution board. Reconnect the drive cables to the rear of the drive holder. The USB cable nests inside the control panel power cable in the right air baffle.

   Ensure that the cable latch clips snap into place on the connectors.

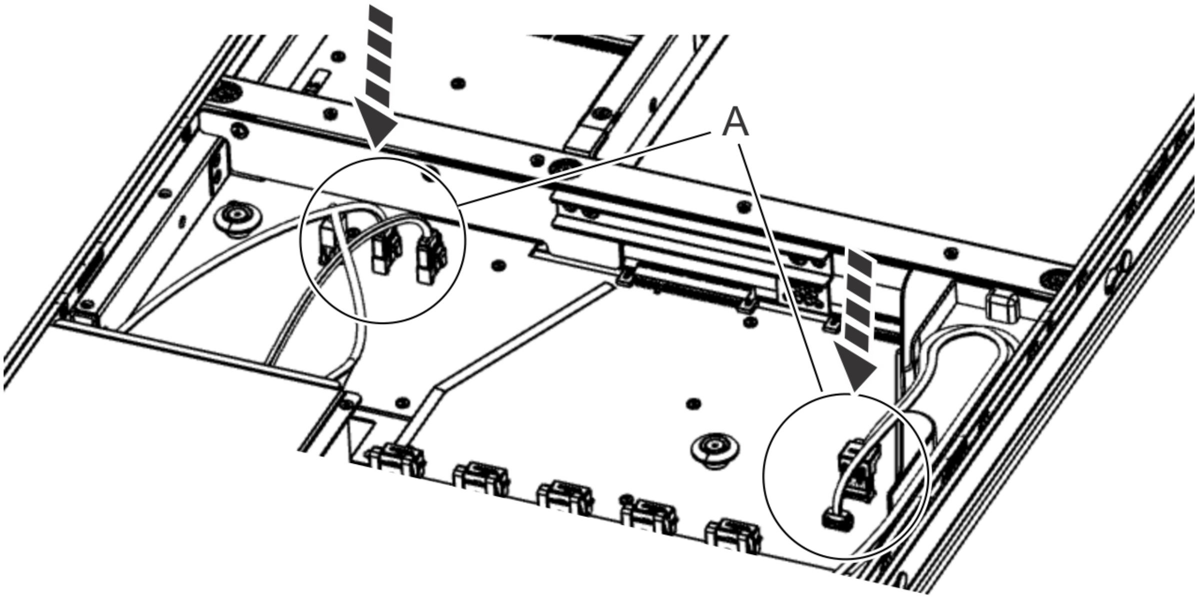

Figure 66. Replacing cables from the power distribution board

7. Install the service access cover.

   For instructions, see "Installing the service access cover on a 7063-CR2 system" on page 105.

8. Replace the system in the rack and replace the components that you removed.

   For instructions, see "Placing a 7063-CR2 system into the operating position" on page 103.

9. Replace the system backplane into the rear of the system.

   a) Ensure that the two system backplane levers are open.

   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

      Important:

      - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
      - Ensure that the system backplane is fully seated and is all the way into the system.
      - You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

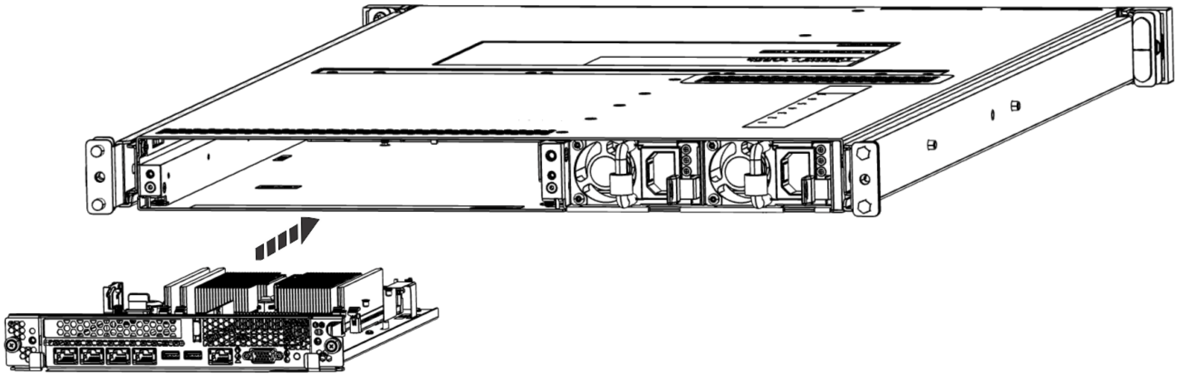

Figure 67. Replacing the system backplane

   c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

    d) Tighten the two screws on the sides of the system backplane.

    e) Using your labels, replace the signal cables into the rear of the system.

    f) Using your labels, replace the two power cords at the rear of the system.

       For instructions, see “Connecting the power cords to a 7063-CR2 system” on page 101.

10. Replace all the fans into the system.

    For instructions, see “Replacing a fan in the 7063-CR2 system” on page 20.

11. Power on the system for operation.

    For instructions, see “Starting the 7063-CR2 system ” on page 98.

Removing and replacing a power supply in the 7063-CR2

Learn how to remove and replace power supplies in the IBM Power Systems HMC (7063-CR2) systems.

Removing a power supply from the 7063-CR2 system

To remove a power supply from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

If a single power supply failed, it can be replaced while the system is running.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

2. Label and remove the power cord from the power supply that you want to remove.

   For instructions, see “Disconnecting the power cords from a 7063-CR2 system” on page 100.

3. To remove the power supply from the system, complete the following steps:

   a) To unseat the power supply from its position in the system, push the locking-tab **(A)** to the left as shown in the following figure.

   b) Grasp the power supply handle with one hand, and pull the power supply partially out of the system.

   c) Place your other hand underneath the power supply and pull the power supply out of the system and place it on an ESD mat.

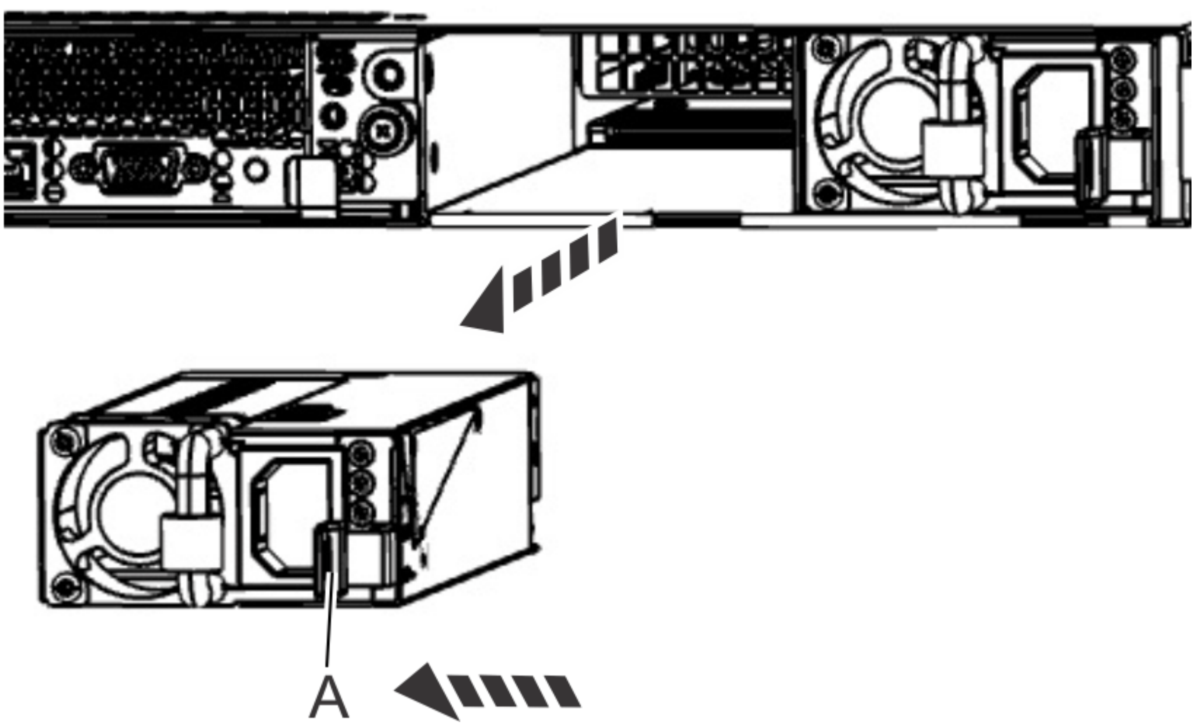

Figure 68. Removing a power supply from the system

Replacing a power supply in the 7063-CR2 system

To replace a power supply in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Wait 30 seconds after you remove a power supply before you install a power supply.
3. To install a power supply in the system, complete the following steps:
   a) Align the power supply with the bay as shown in the following figure. The fan is on the left; the plug is on the right.
   b) Slide the power supply into the system until the latch locks in place.

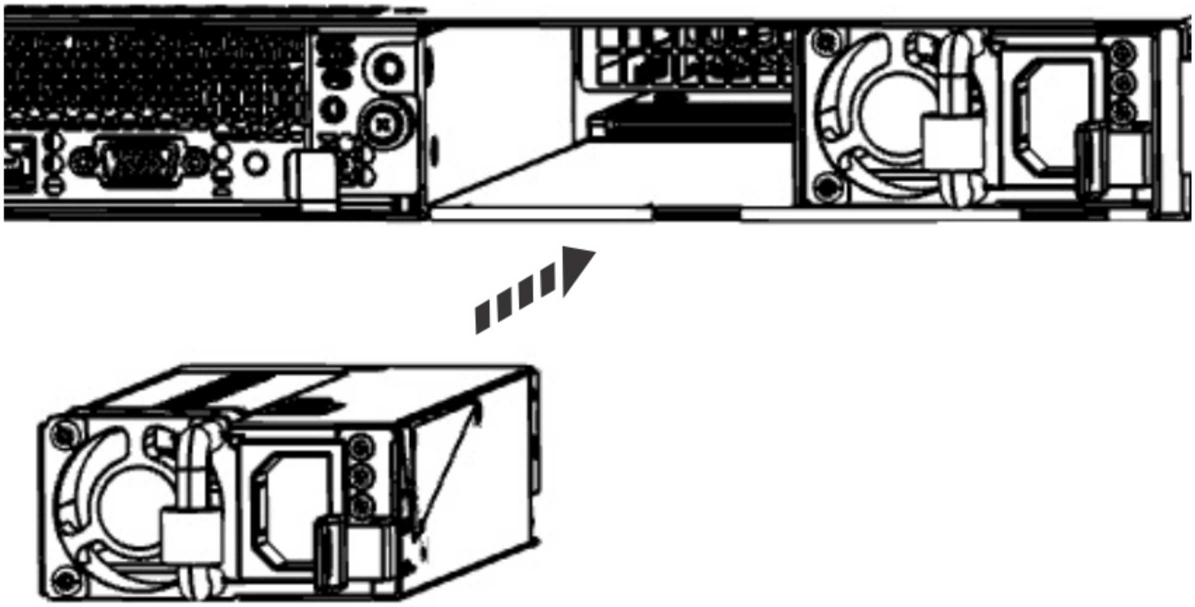

Figure 69. Installing a power supply in the system

4. Reconnect the power cord.

   For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

Removing and replacing the system backplane in the 7063-CR2

Learn how to remove and replace the system backplane in the IBM Power Systems HMC (7063-CR2) system.

Before you begin

Removing or replacing this part is a customer task. You can complete this task yourself, or contact a service provider to complete the task for you. You might be charged a fee by the service provider for this service.

Before you begin replacing the system backplane, write down the system serial number and machine model type. After you replace the system backplane, you must set the system serial number and machine model type in the system backplane.

Preparing the 7063-CR2 system to remove the system backplane

To prepare to remove the system backplane from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Make note of the system serial number and the machine model type. After you replace the system backplane, you must set the system serial number and machine model type in the system backplane.
2. Make note of the BMC network settings. Record the BMC IP settings.

   After you replace the system backplane, you might need to reconfigure the BMC network settings.

Removing the system backplane from the 7063-CR2 system

To remove the system backplane from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

You can use a magnetic tip screwdriver to remove and replace the screws.

As part of the system backplane replacement, the system processor modules are moved from the old system backplane to the new system backplane.

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.
2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.
3. Remove the system backplane from the rear of the system.

   a) Label and remove the two power cables.

   For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.

   b) Label and remove the signal cables from the rear of the system.

   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

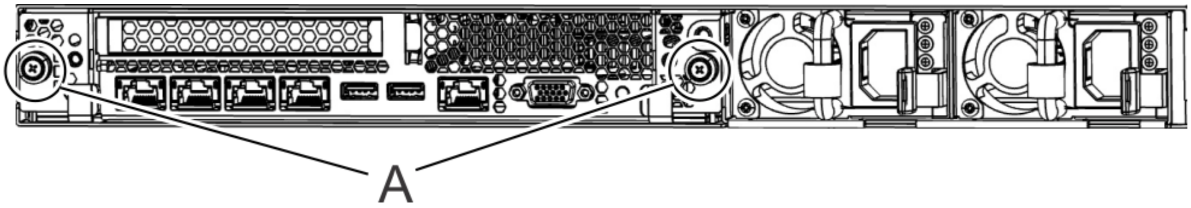

Figure 70. Removing the system backplane screws

   d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

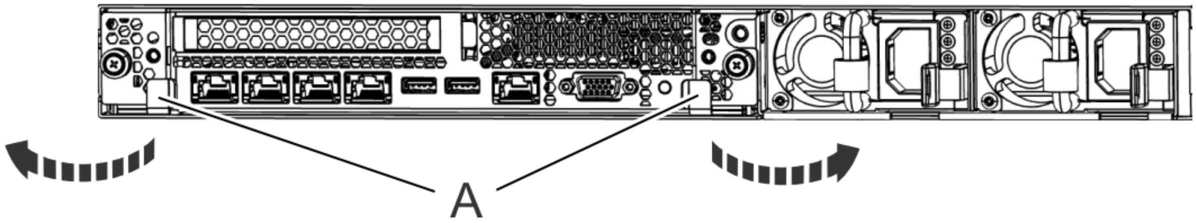

Figure 71. Unlatching the system backplane

e) Support the system backplane by the bottom as you slide it from the system.

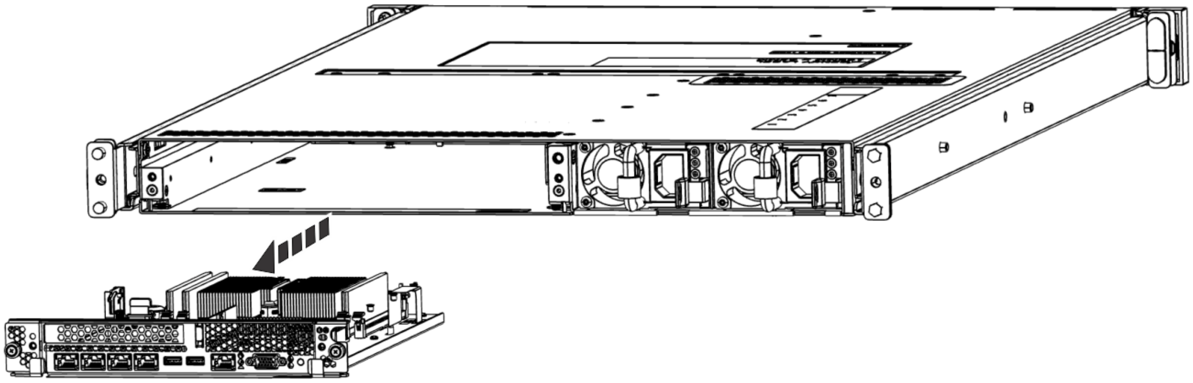

Figure 72. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

Replacing the system backplane in the 7063-CR2 system

To replace the system backplane in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Remove the replacement system backplane from the static-protective package and place it on an ESD mat next to the old system backplane.

The following steps move the system processor module from the old system backplane to the new system backplane:

3. Loosen the load arm screw **(A)** of the system processor heat sink **(B)** that you are removing with a T20 hexalobular driver. The load arm pivots up in the direction that is shown in the following figure.

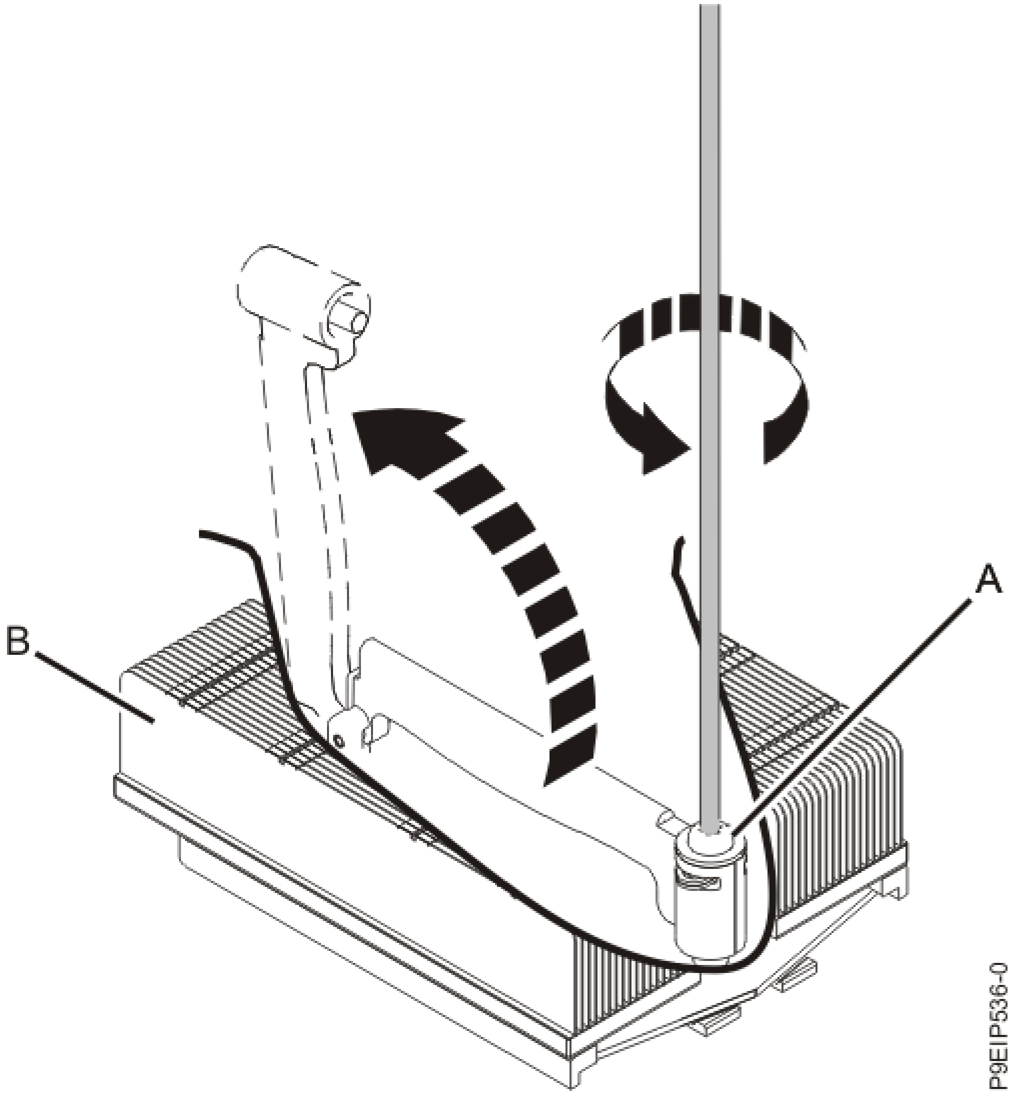

Figure 73. Loosening the load arm screw of the heat sink

4. Grip the heat sink and remove it by lifting it straight up as shown in the following figure.

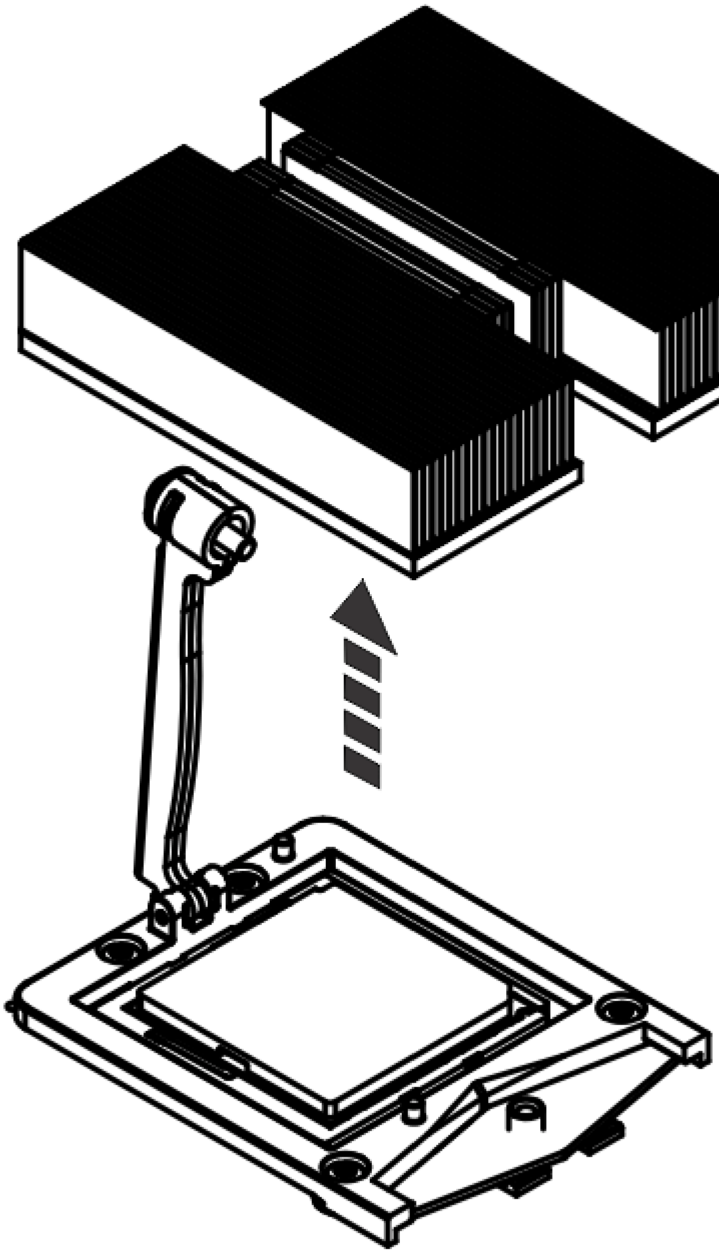

Figure 74. Removing the heat sink

5. Place the heat sink upside down on a clean surface.
6. Using tweezers, carefully remove the TIM from the top of the system processor module and place it in a clean, dry area.

   The TIM can tear easily.
7. Remove the cover from the system processor socket on the new system backplane
8. Inspect the system processor socket area and remove any dust or debris (use a can of compressed air).
9. Align the tool with the beveled edge **(A)** of the system processor module as shown in the following figure. Lower the tool over the system processor module by ensuring the two guide pins **(C)** are inserted into the alignment holes **(B)** on each side of the tool.

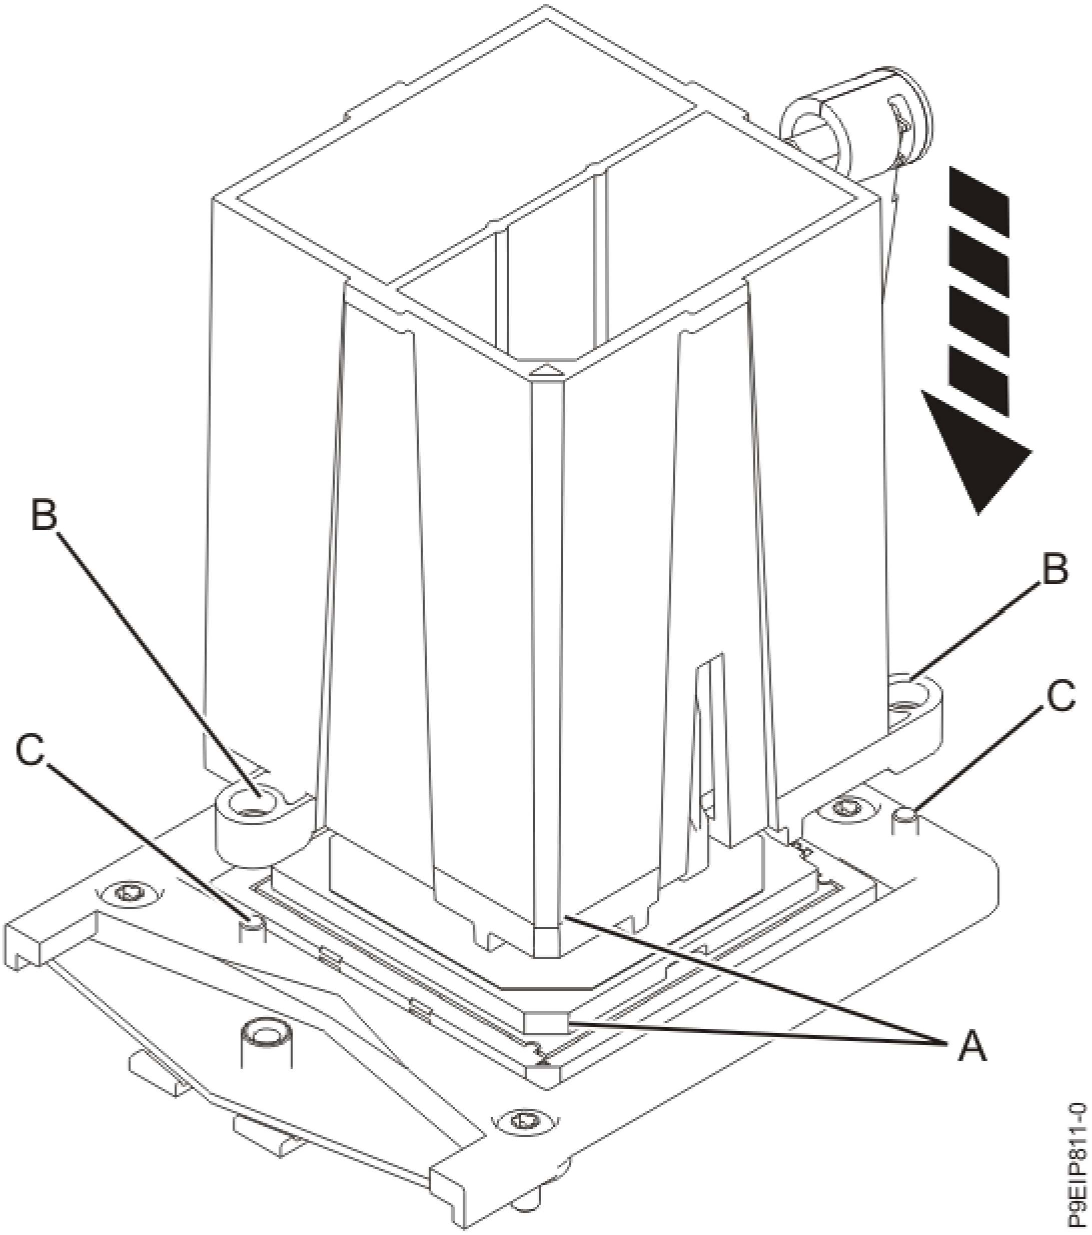

Figure 75. Lowering the removal tool onto the system processor module

10. Using the lift tool, move the system processor module from the old system backplane socket to transfer it to the new system backplane socket.
11. With the removal tool **(A)** sitting on top of the system processor module, push down on the tool to lock the system processor module into the tool as shown in the following figure.

    The tool drops slightly when you push down on the system processor module so that the jaws can grab the bottom of the module. Make sure that both of the tool jaws are locked on the system processor module. Do not press the blue release tabs until directed to do so later.

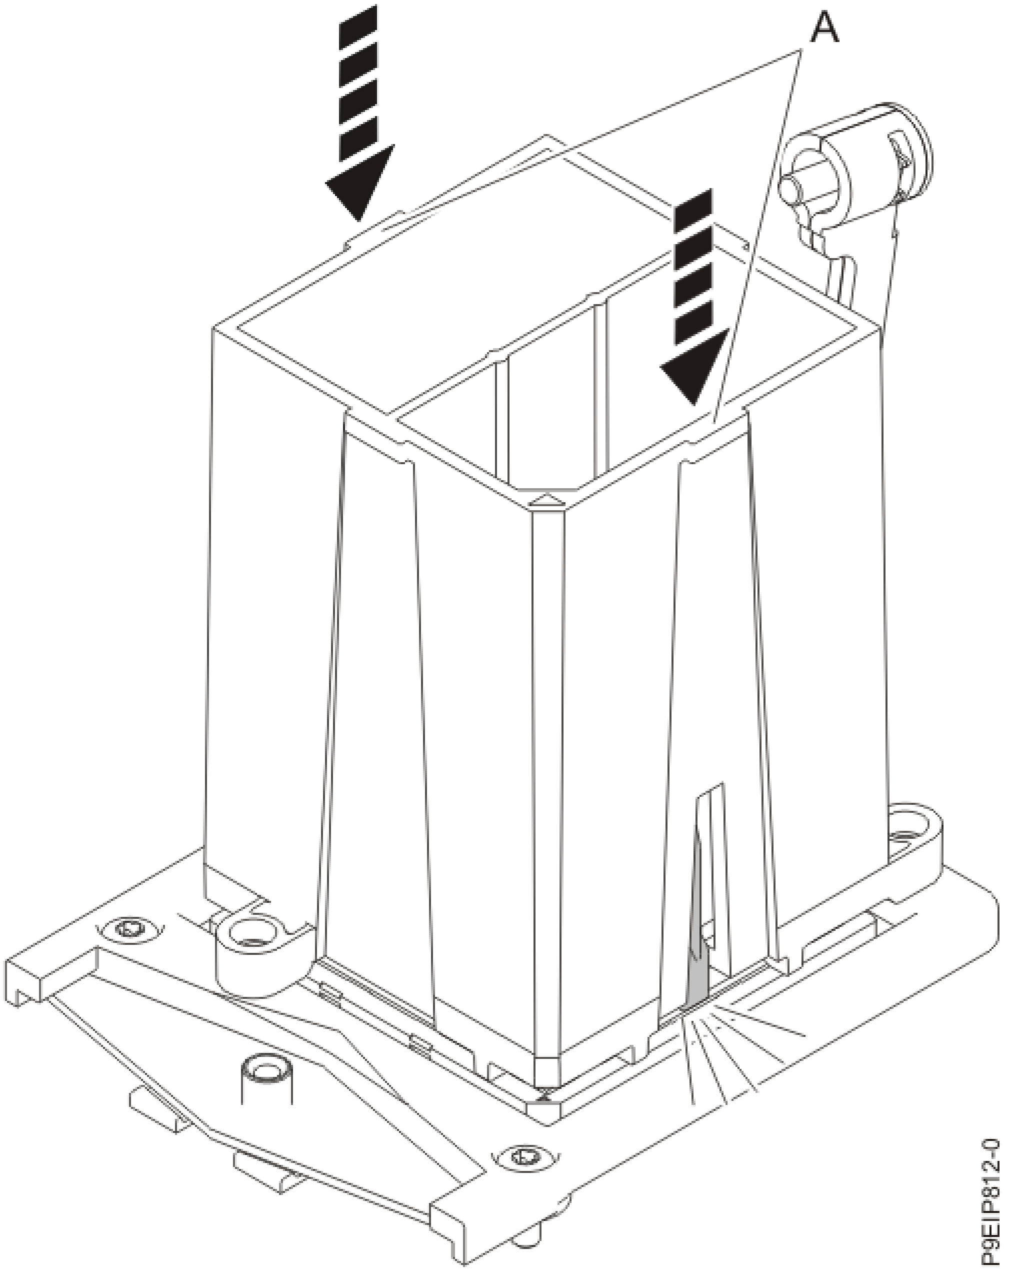

Figure 76. Locking the system processor module into the tool

12. Lower the tool and system processor module onto the socket. Align the beveled corner **(A)** of the tool with the beveled corner on the socket as shown in the following figure.

    Ensure that the two guide pins **(C)** are inserted into the alignment holes **(B)** on each side of the tool. Use care to lower the tool evenly without tilting the tool. Do not attempt to slide the tool and the system processor module in any direction while the system processor module is touching the socket. If the tool and the system processor module are not aligned with the guide pins, lift the tool and the system processor module and reposition them.

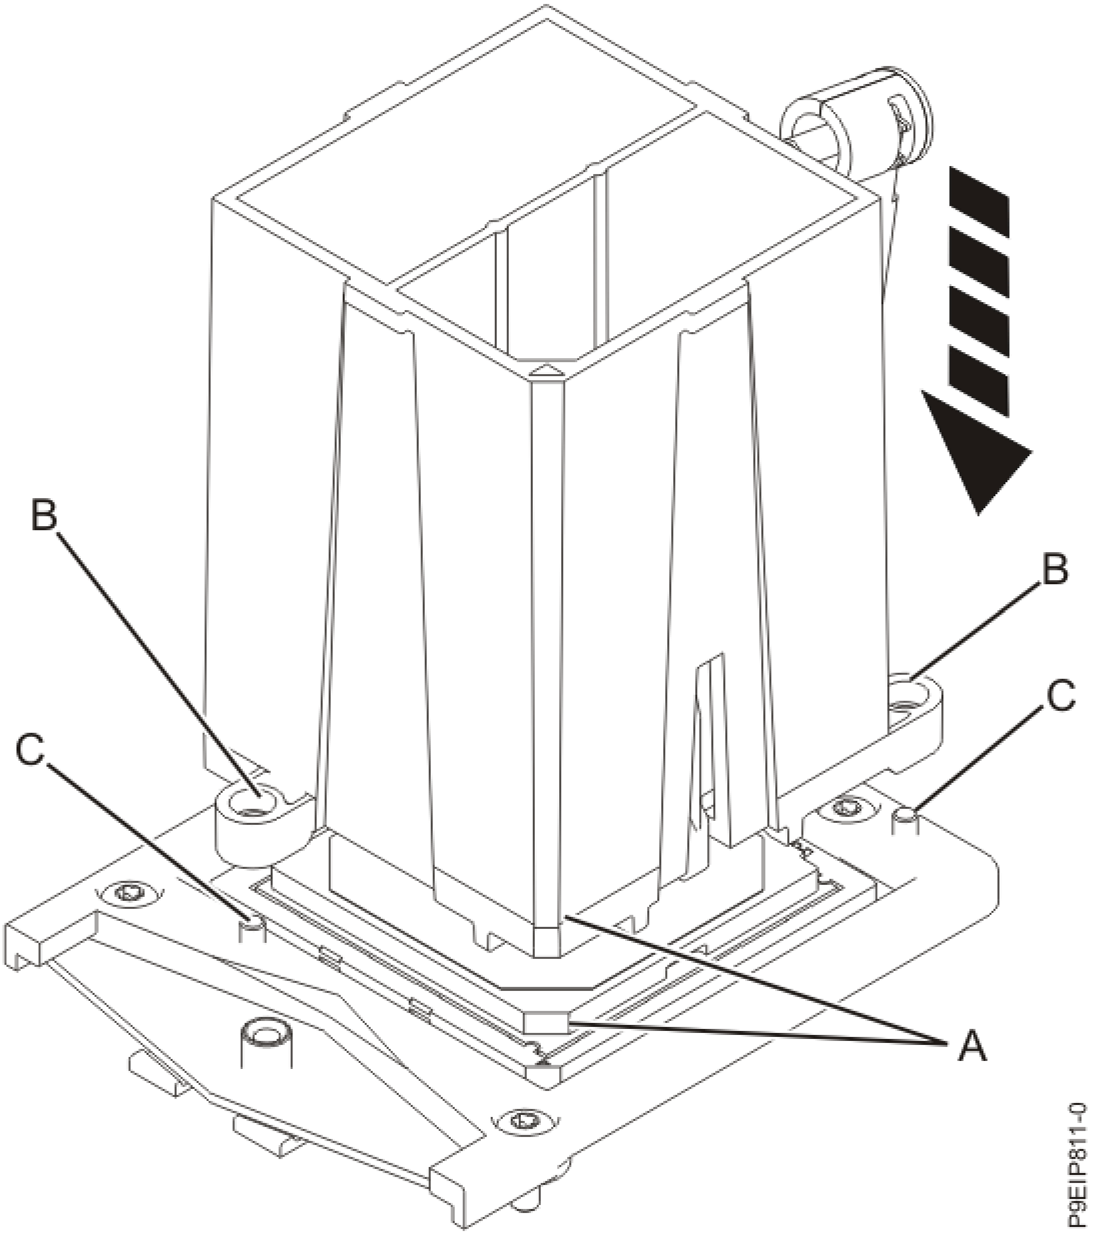

Figure 77. Installing the system processor module

13. After the tool and system processor module holes and guide pins are properly aligned, squeeze and hold the two blue release tabs **(A)** together until a firm stop is reached as shown in the following figure.

    Then, lift the tool off the system processor module.

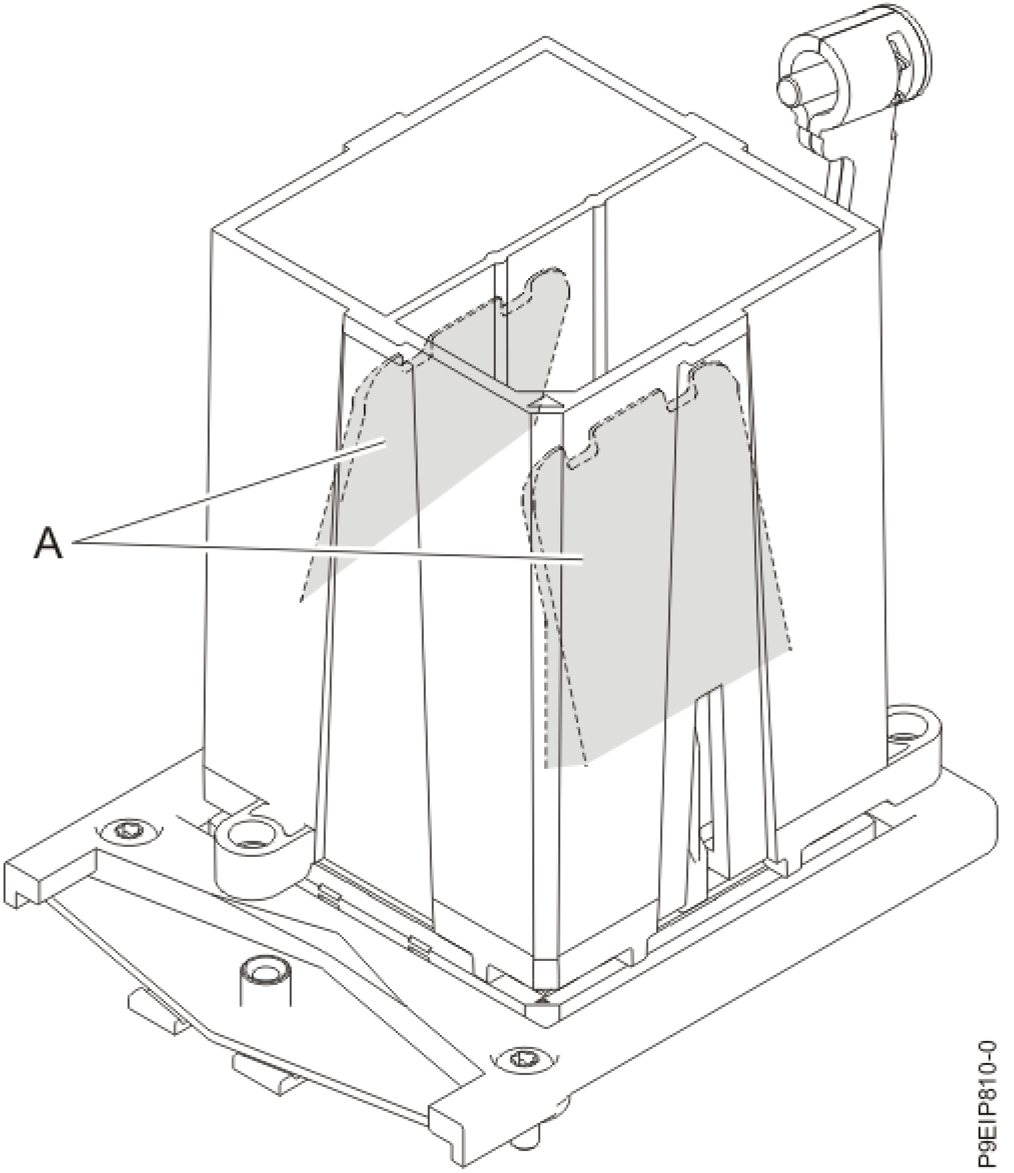

Figure 78. Removing the system processor module tool

14. Inspect the thermal interface material (TIM) for visible signs of damage. If you see folds, tears, bends, or if you have doubts about the TIM, replace it.

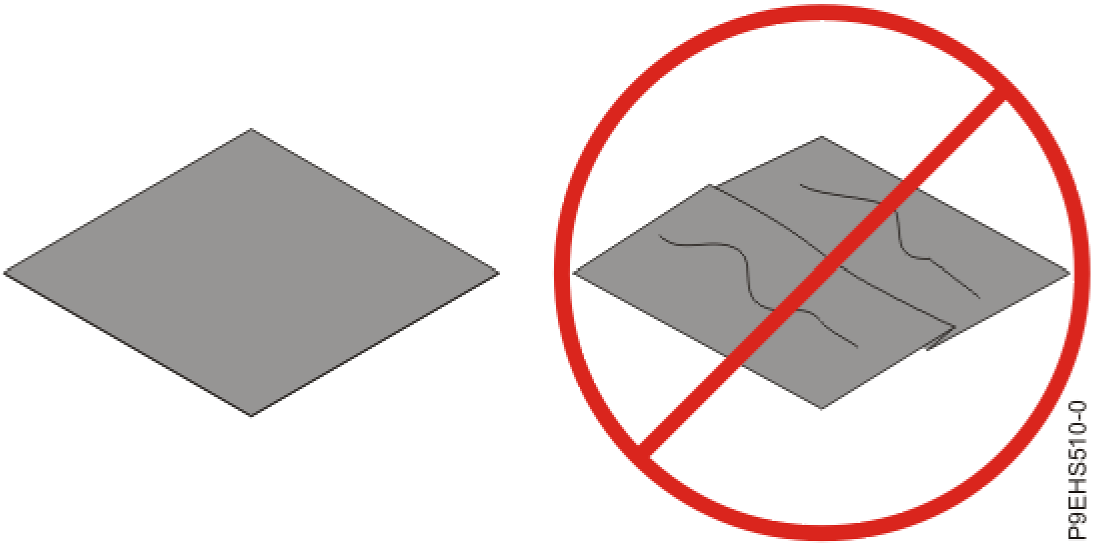

Figure 79. Inspecting the thermal interface material

15. Choose one of the following repair options:

Option	Description
Is the TIM damaged?	It is damaged. Proceed to step “16” on page 62 to replace the TIM and install the existing heat sink.
Is the TIM OK?	It is not damaged and can be reused. Proceed to step “18” on page 62 to reuse the TIM and install the existing heat sink.

16. Use this step to install a new TIM and reuse the existing heat sink.
    a) Open the TIM packaging and carefully remove the TIM, holding it by the edges of the carrier strip and holding it away from the shipping container.
    b) Remove the protective film from the clear carrier strip by using the supplied tweezers.

       Note: The TIM must remain flat. Small wrinkles are acceptable, but folds are not acceptable.
    c) Using the tweezers, remove the TIM from the carrier strip and center it onto the system processor module.

       The TIM has no preferred up side. The TIM can be placed on the system processor module and centered..
17. Continue with step “19” on page 62.
18. Use this step to reuse the existing undamaged TIM and heat sink.
    a) Using the tweezers, move the old TIM from the clean, dry surface and center it onto the new system processor module.

       The TIM has no preferred up side. The TIM can be placed on the system processor module and centered.
19. Carefully lower the heat sink over the system processor module, ensuring that the holes in the heat sink align with the two guide pins **(A)** on the socket, as shown in the following figure.

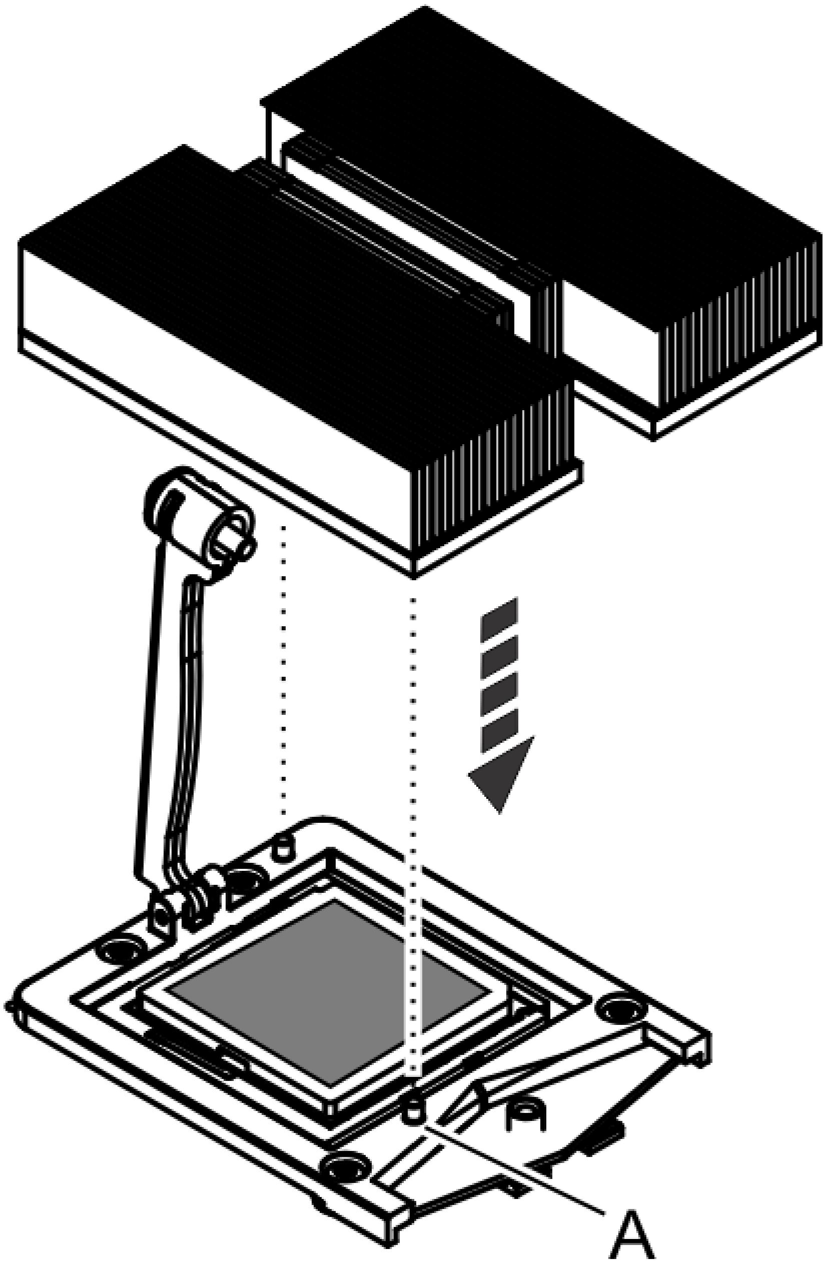

Figure 80. Installing the heat sink

20. Move the load arm **(A)** into position over the heat sink **(B)** and tighten the load arm screw with a T20 hexalobular driver, as shown in the following figure.

    Note: Do not over tighten the load arm screw.

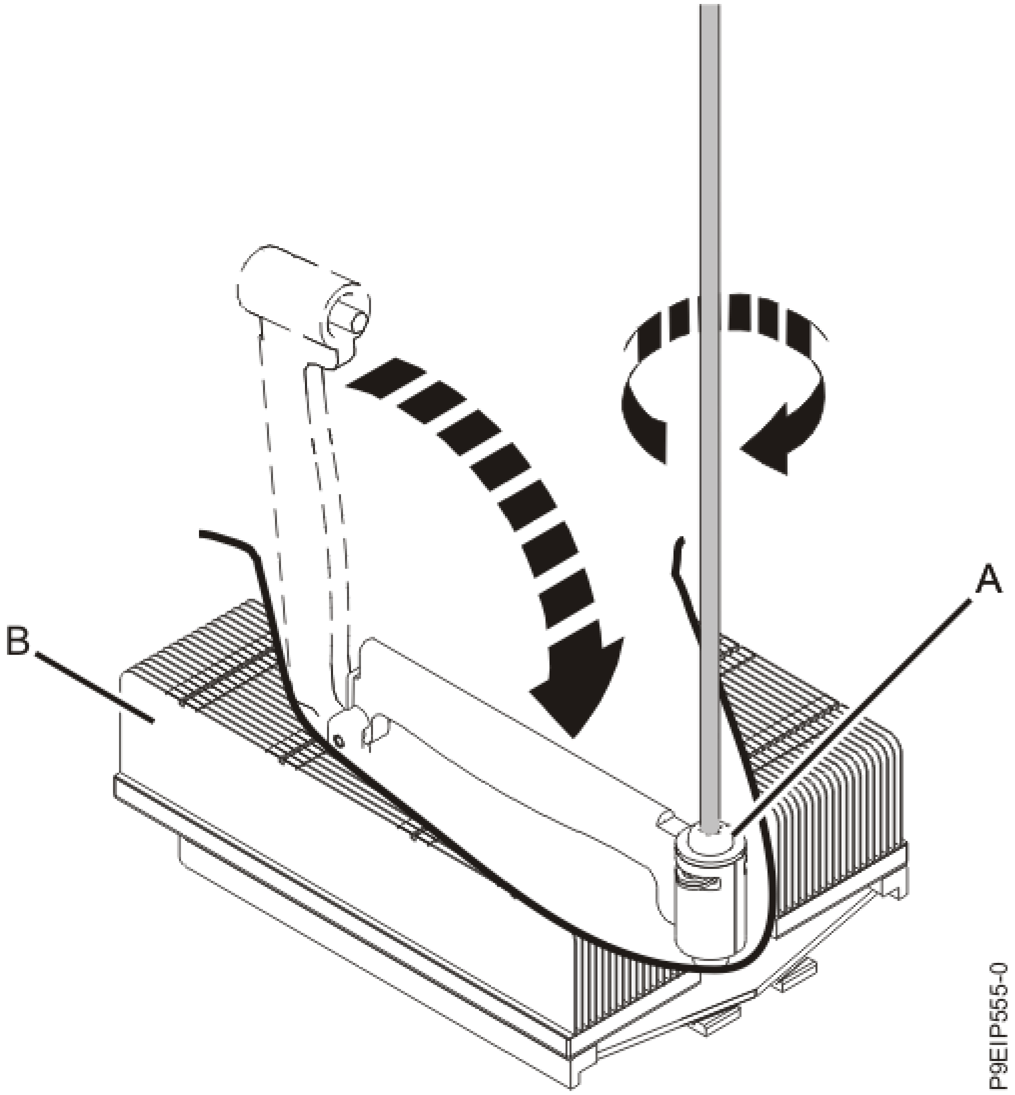

Figure 81. Tightening the load arm screw

The following steps move the remaining parts from the old system backplane to the new system backplane:

21. Move the memory DIMMs from the old system backplane to the corresponding location on the new system backplane.

    For instructions, see "Removing and replacing memory in the 7063-CR2" on page 21.

22. Move the TPM card from the old system backplane to the corresponding location on the new system backplane.

    For instructions, see "Removing and replacing the trusted platform module in the 7063-CR2" on page 89.

23. If applicable, move the PCIe adapter from the old system backplane to the corresponding location on the new system backplane.

    a) Remove the screw that secures the PCIe adapter to the system backplane.

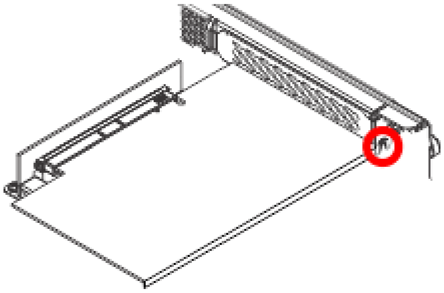

Figure 82. Removing the tail-stock screw

b) Open the retainer clip that secures the PCIe adapter to the PCIe riser, by moving the blue lever to the unlocked position.

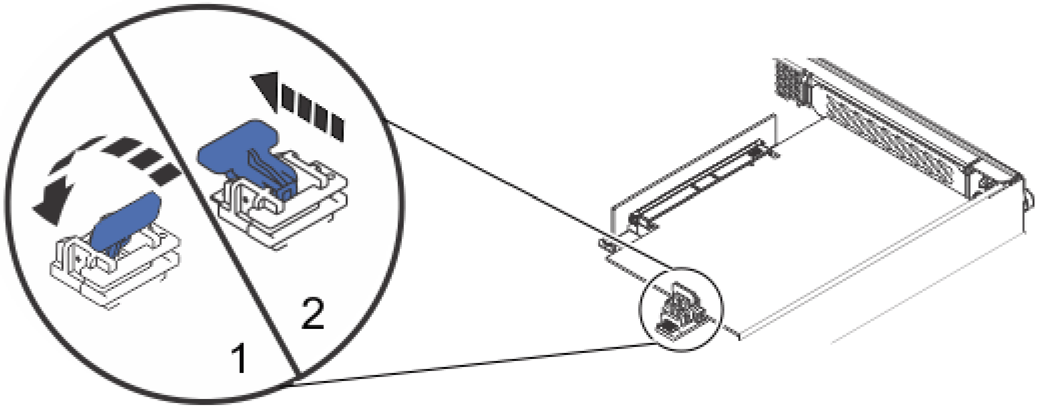

Figure 83. Removing the PCIe adapter retaining latch

c) Move the retainer clip away from the PCIe adapter.

d) Remove the PCIe adapter from the PCIe riser on the old system backplane.

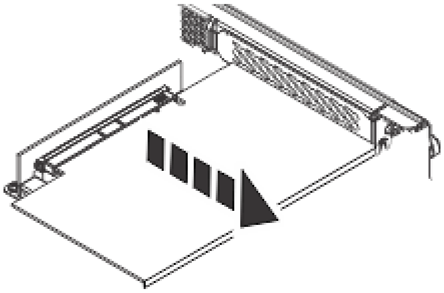

Figure 84. Removing the PCIe adapter

e) Replace the PCIe adapter into the PCIe riser on the new system backplane

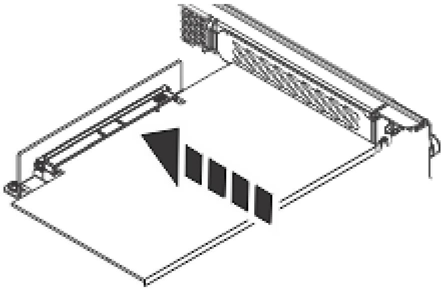

Figure 85. Replacing the PCIe adapter

f) Replace the screw that secures the PCIe adapter to the system backplane.

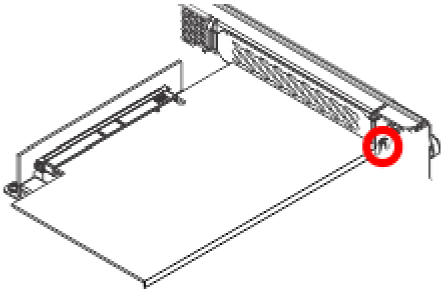

Figure 86. Replacing the tail-stock screw

g) Move the retainer clip to secure the PCIe adapter to the PCIe riser. Ensure that the clip is fully seated around the edge of the adapter.

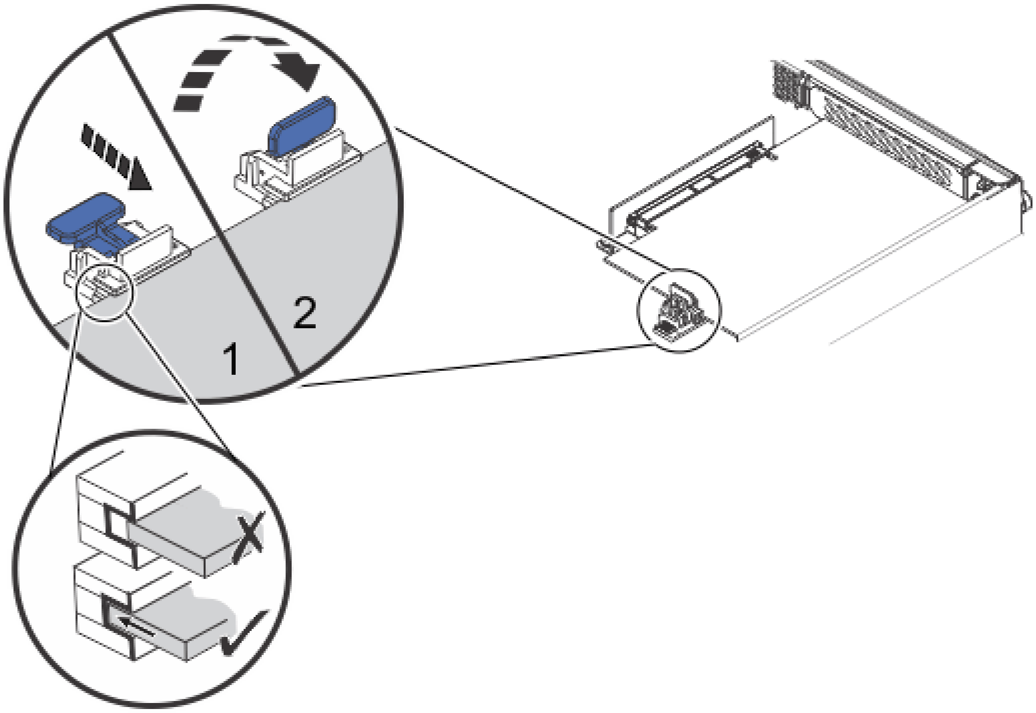

Figure 87. Replacing the PCIe adapter retaining latch

h) Close the retainer clip that secures the PCIe adapter to the PCIe riser.

24. Remove the blue protective insulator from the battery.
25. Replace the system backplane into the rear of the system.

a) Ensure that the two system backplane levers are open.

b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

Important:

- Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
- Ensure that the system backplane is fully seated and is all the way into the system.
- You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

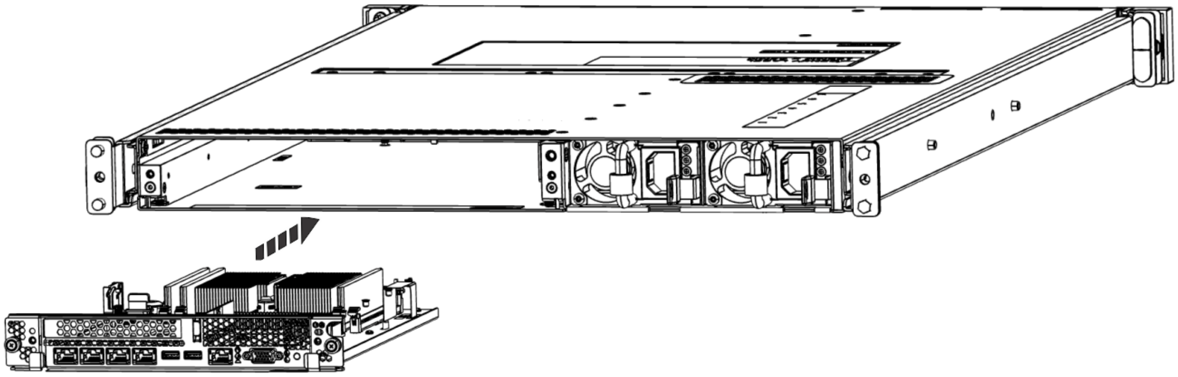

Figure 88. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

Preparing the 7063-CR2 system for operation after removing and replacing the system backplane

To prepare the IBM Power Systems HMC (7063-CR2) system for operation after removing and replacing the system backplane, complete the steps in this procedure.

Procedure

1. The serial number needs to be set before the operating system starts to avoid issues.

   Follow the steps listed here: Scale-out LC system VPD update tool (www14.software.ibm.com/webapp/set2/sas/f/lopdiags/scaleOutLCdebugtool.html#OpenPOWER).
2. Ensure that the BMC network settings are correct. For instructions, see Configuring the BMC IP address.
3. The date and time needs to be set before the operating system starts to avoid issues.

   Follow these petitboot steps:

   a) From the petitboot menu, select **Exit to shell**.

   b) Run the following two commands to check the date and time:

   ```
   date
   ipmitool sel time get
   ```

   c) To set the correct date and time in UTC use the following format:

   ```
   date -s YYYY.MM.DD-HH:MM
   ```

   d) Run the following ipmitool command to set the system entry log (SEL) date and time to the new values:

   ```
   ipmitool sel time set now
   ```

   e) Verify that the ipmitool and date commands now show the correct values by running the following commands again:

   ```
   date
   ipmitool sel time get
   ```

   f) To leave the petitboot shell, type **exit**.

4. The operating system of the HMC needs to have valid credentials with which to access the BMC.

   The credentials are set during the setup of the HMC. The credentials can also be set by running a task. If the credentials are not set up, the HMC cannot run the call-home functions for itself. Reconfigure the credentials after you sign in to the HMC.

   To run the task, click the **HMC Management** icon, then select **Console Settings** > **Console Inband Communication Credentials**
5. If secure boot option is enabled before the replacement of the system backplane, then the user needs to re-enable the secure boot option by following the steps in Enabling secure boot on 7063-CR2 HMC.

Removing and replacing the system processor module in the 7063-CR2

Learn how to remove and replace the system processor module in the IBM Power Systems HMC (7063-CR2) system.

Before you begin

Removing or replacing this part is a customer task. You can complete this task yourself, or contact a service provider to complete the task for you. You might be charged a fee by the service provider for this service.

Removing the system processor module from the 7063-CR2 system

To remove the system processor module from the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

(L007)

CAUTION: A hot surface nearby. (L007)

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.
2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   Attention:

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.

- If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

3. Remove the system backplane from the rear of the system.
   a) Label and remove the two power cables.

      For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.
   b) Label and remove the signal cables from the rear of the system.
   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

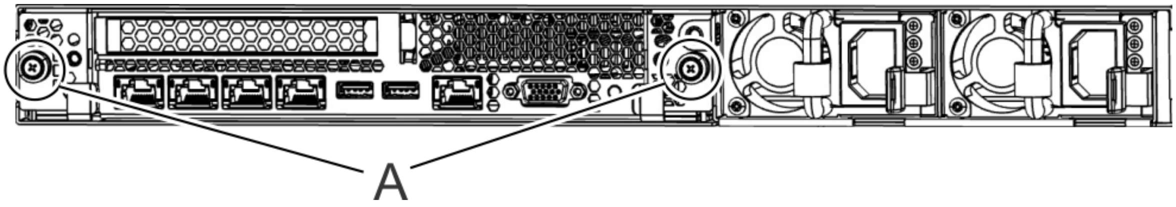

Figure 89. Removing the system backplane screws

   d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

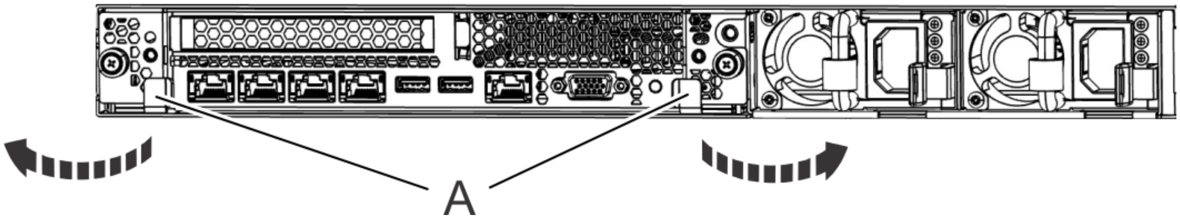

Figure 90. Unlatching the system backplane

   e) Support the system backplane by the bottom as you slide it from the system.

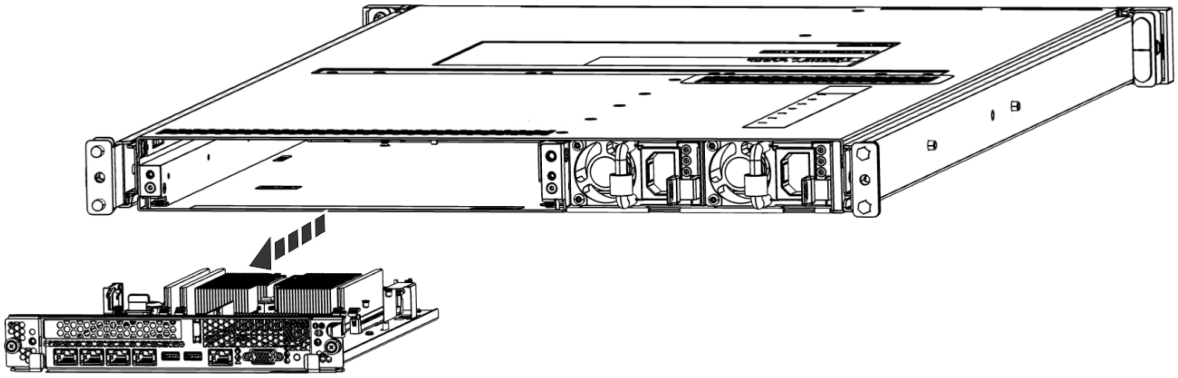

Figure 91. Removing the system backplane

   f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

4. Open the packaging of the new system processor module and place the cover upside down next to the tray, as shown in the following figure. The cover is used for the system processor module that you want to replace.

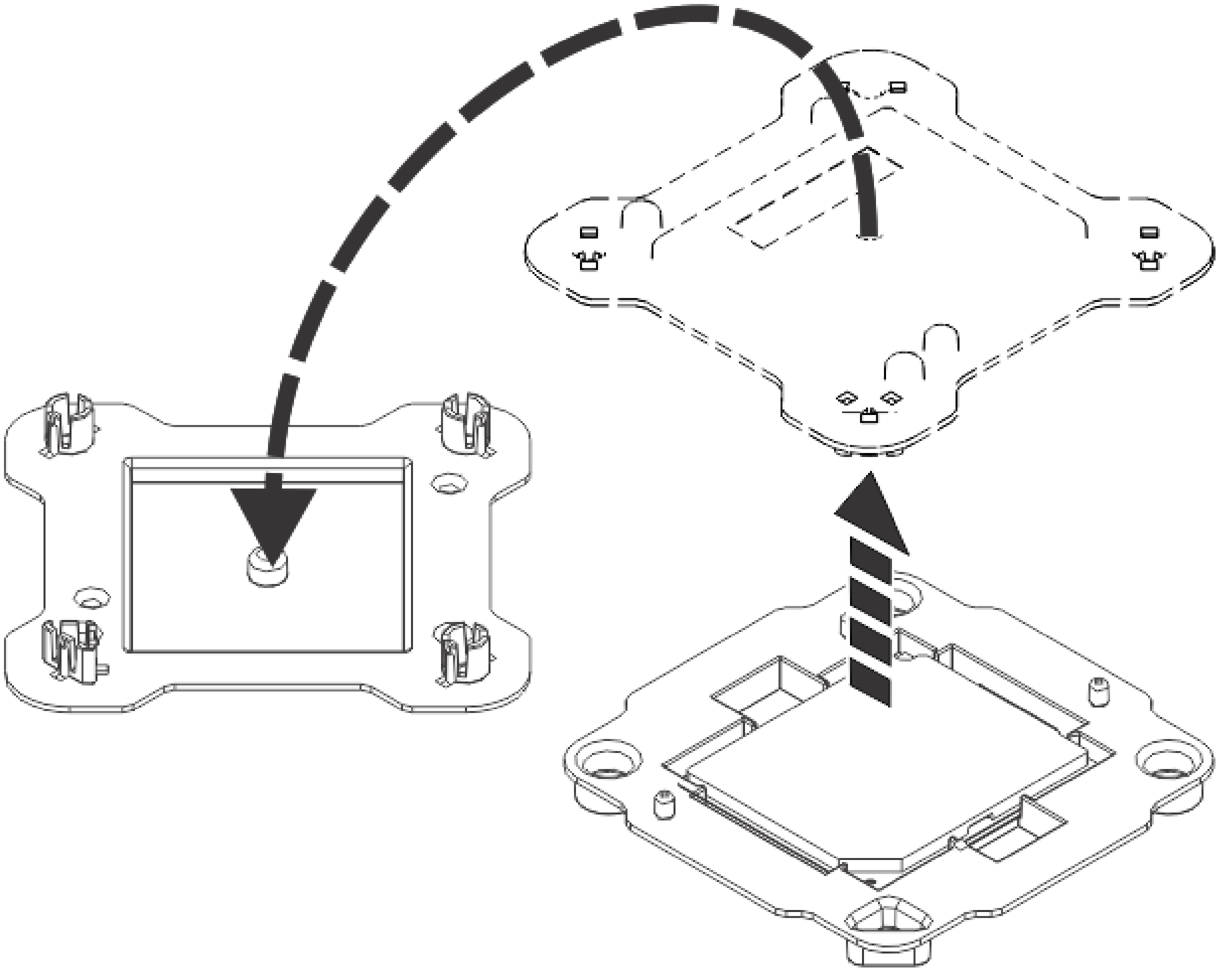

Figure 92. Opening the system processor module packaging

5. Loosen the load arm screw **(A)** of the system processor heat sink **(B)** that you are removing with a T20 hexalobular driver. The load arm pivots up in the direction that is shown in the following figure.

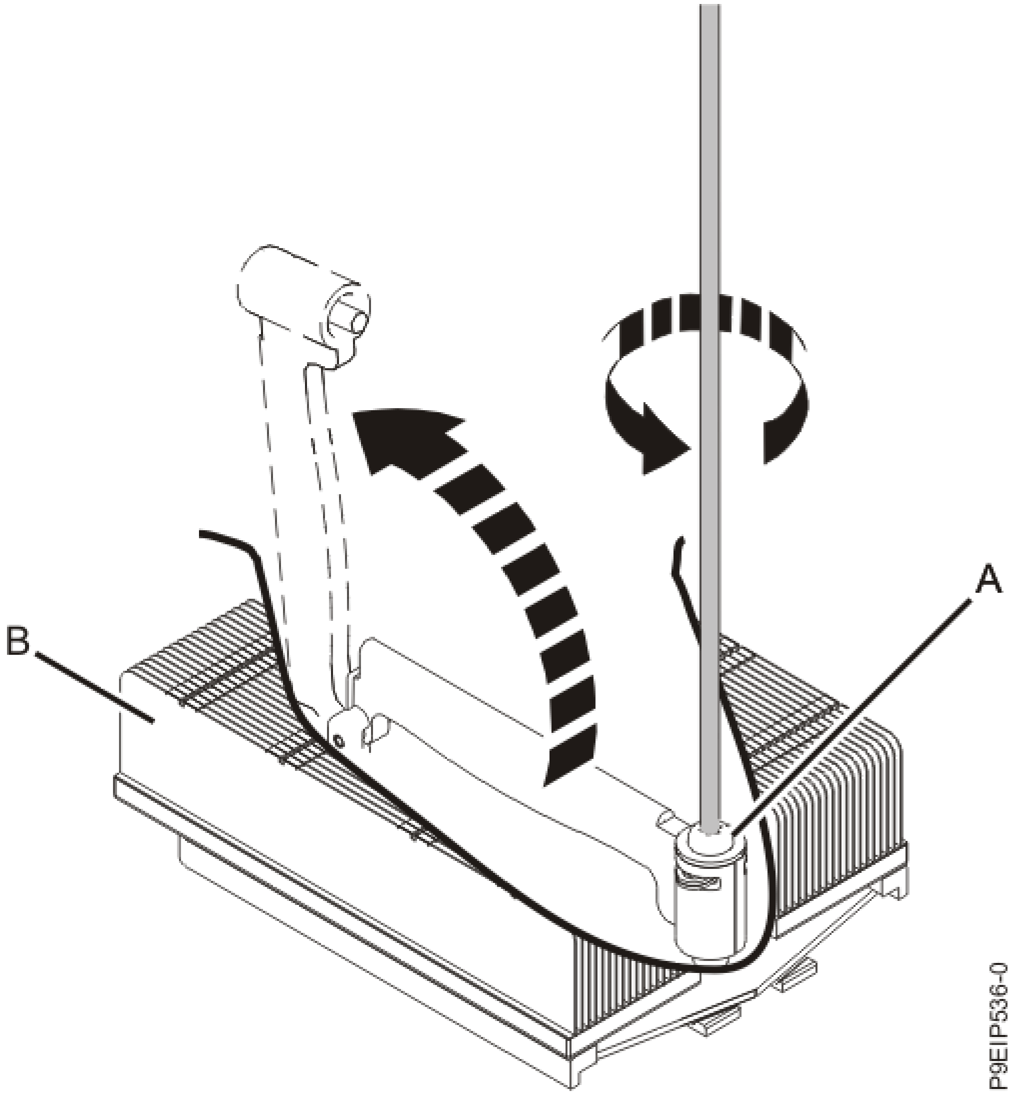

Figure 93. Loosening the load arm screw of the heat sink

6. Grip the heat sink and remove it by lifting it straight up as shown in the following figure.

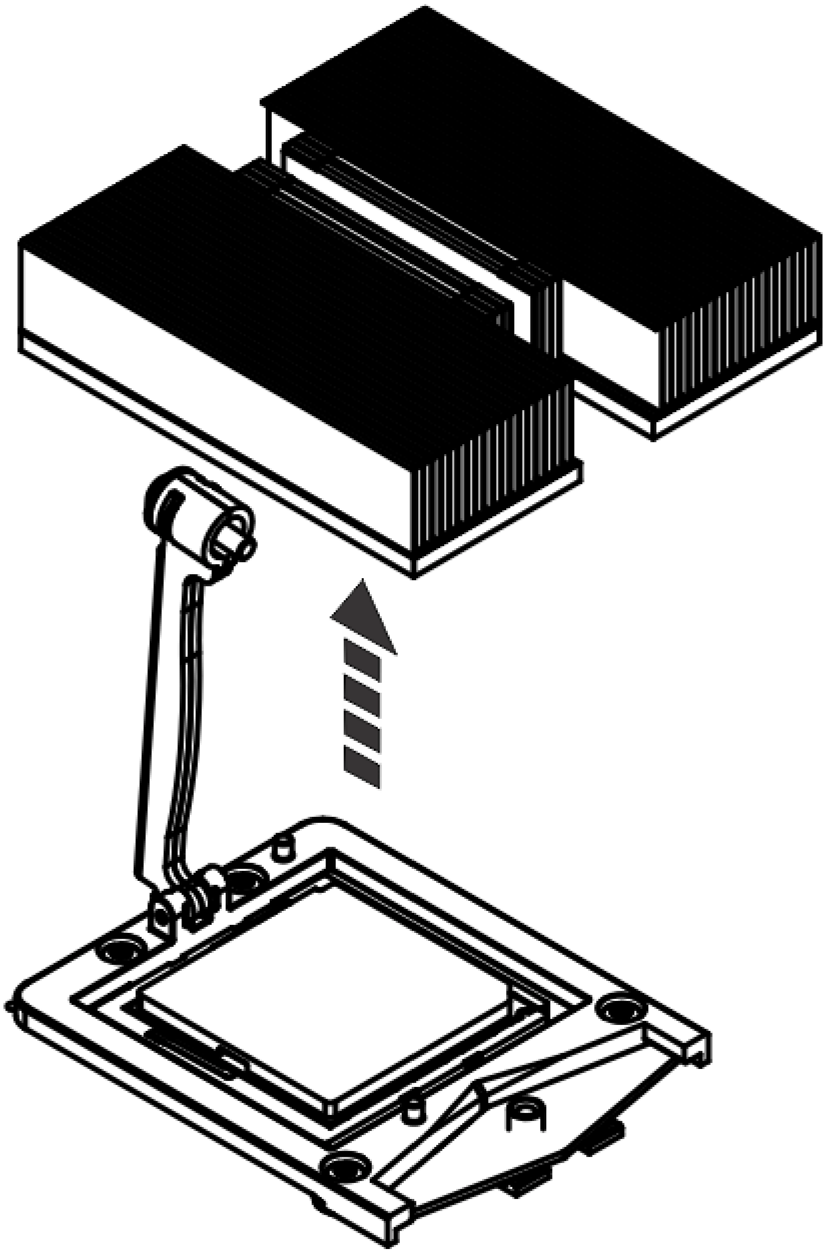

Figure 94. Removing the heat sink

7. Place the heat sink upside down on a clean surface.
8. Using tweezers, carefully remove the TIM from the top of the system processor module and place it in a clean, dry area.

   The TIM can tear easily.
9. Inspect the system processor socket area and remove any dust or debris (use a can of compressed air).
10. Align the tool with the beveled edge **(A)** of the system processor module as shown in the following figure. Lower the tool over the system processor module by ensuring the two guide pins **(C)** are inserted into the alignment holes **(B)** on each side of the tool.

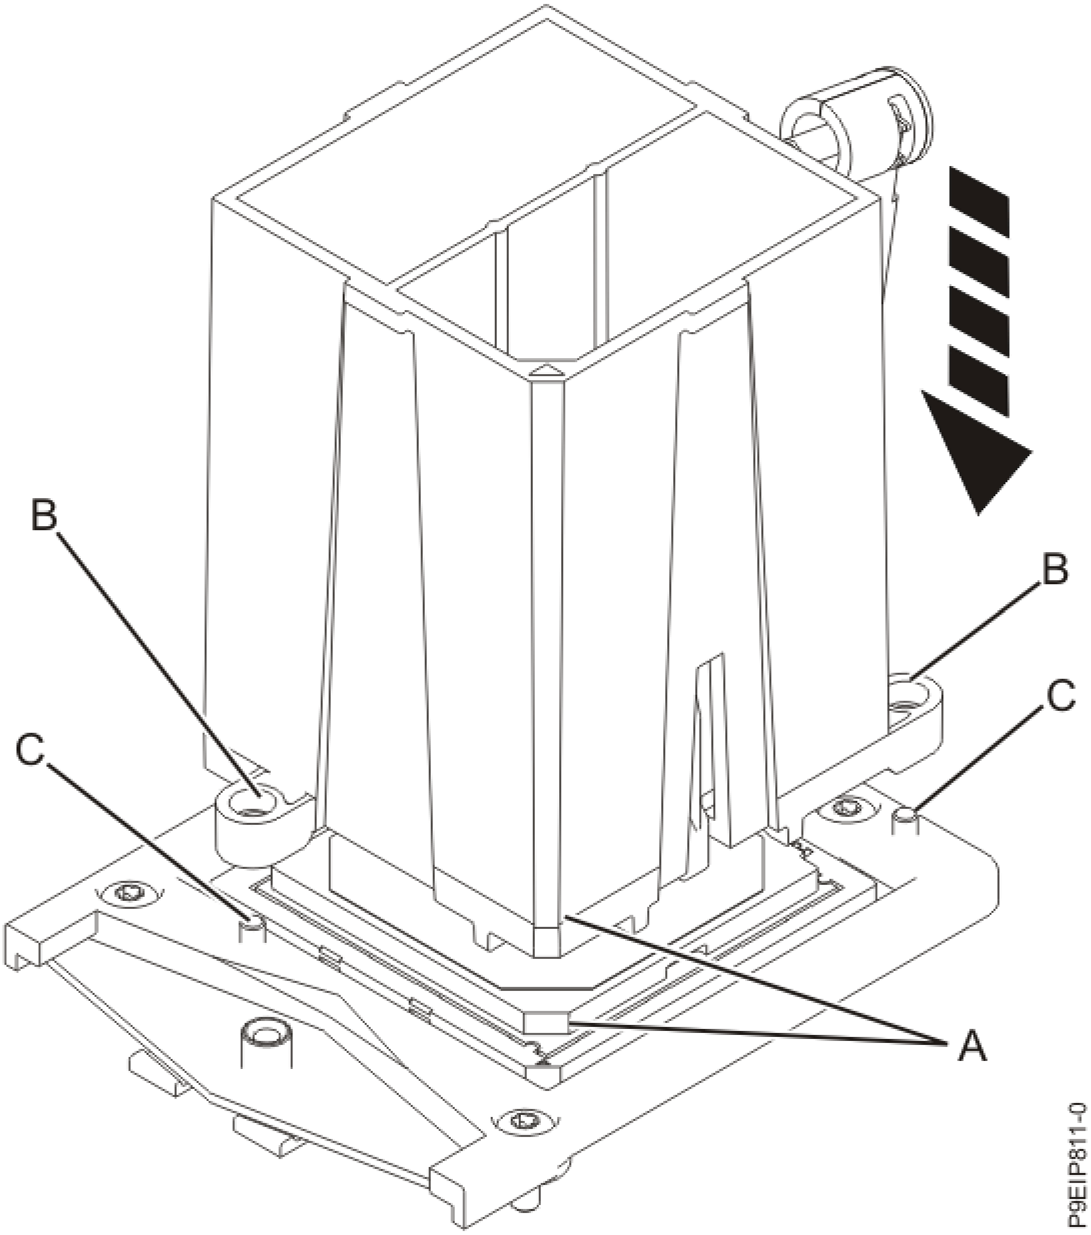

Figure 95. Lowering the removal tool onto the system processor module

11. With the removal tool **(A)** sitting on top of the system processor module, push down on the tool to lock the system processor module into the tool as shown in the following figure.

    The tool drops slightly when you push down on the system processor module so that the jaws can grab the bottom of the module. Make sure that both of the tool jaws are locked on the system processor module. Do not press the blue release tabs until directed to do so later.

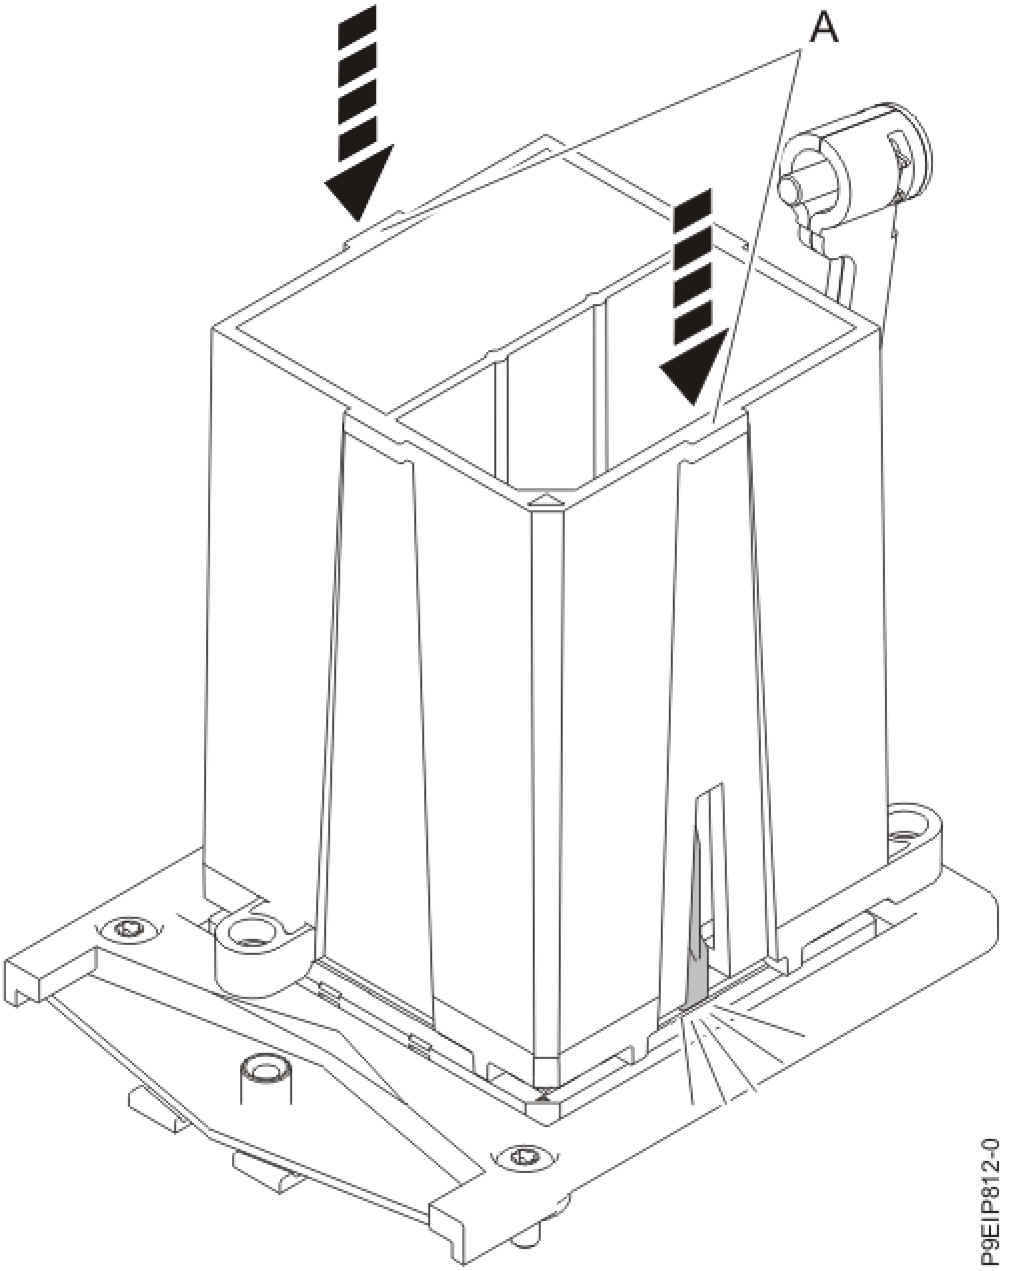

Figure 96. Locking the system processor module into the tool

12. Hold the outside of the tool and lift the tool and system processor module from the socket. Place them at an angle on the top cover of the system processor module packaging as shown in the following figure.

    Setting the system processor module at an angle on the top cover of the system processor module packaging will make it easier to pick up and place in the packaging after you replace the system processor module.

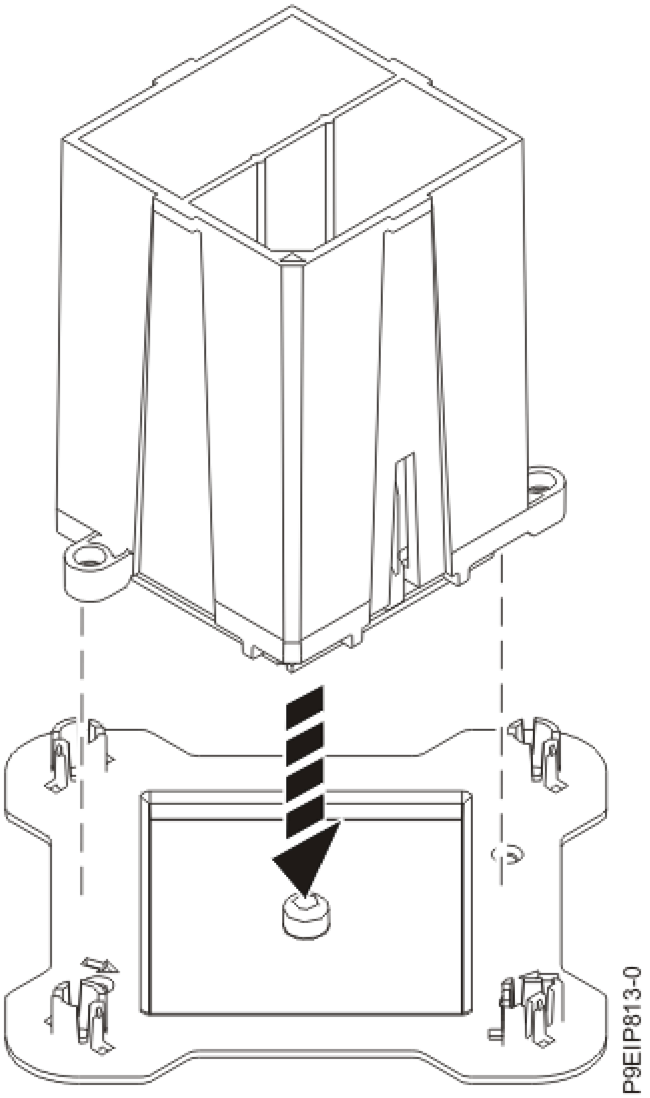

Figure 97. Placing the system processor module at an angle on the top cover of the packaging

13. Squeeze the two blue tabs to release the system processor module from the tool as shown in the following figure.

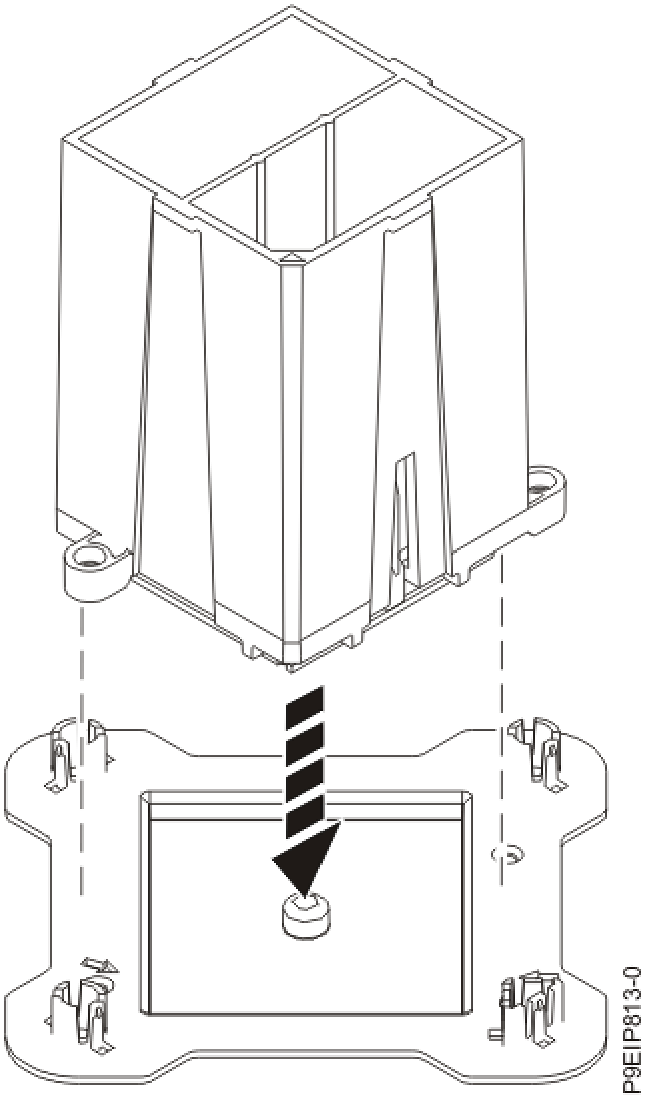

Figure 98. Releasing the system processor module from the tool

Replacing the system processor module in the 7063-CR2 system

To replace the system processor module in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.
2. Inspect the system processor socket area and remove any dust or debris (use a can of compressed air).
3. Remove the replacement processor module from the shipping tray. Align the beveled corner **(A)** of the tool over the beveled corner of the module as shown in the following figure. Ensure that the guide **(B)** fits into the alignment pin **(C)**.

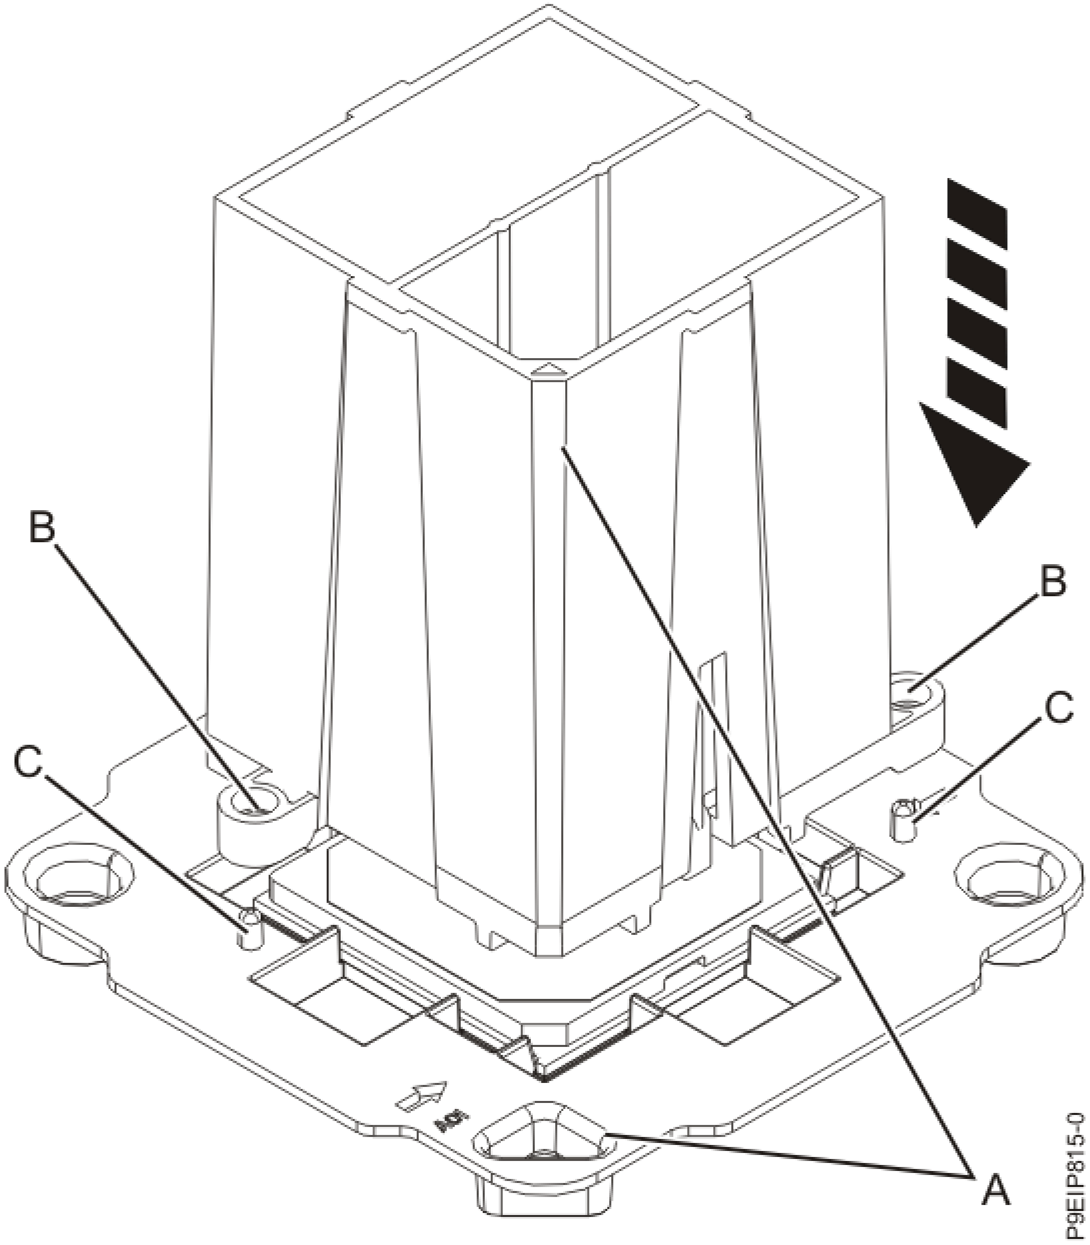

Figure 99. Aligning the removal tool

4. With the removal tool sitting on top of the system processor module, push down on the tool to lock the system processor module into the tool as shown in the following figure.

   The tool drops slightly when you push down on the system processor module so that the jaws can grab the bottom of the module. Make sure that both of the tool jaws are locked on the system processor module. Do not press the blue release tabs until directed to do so later.

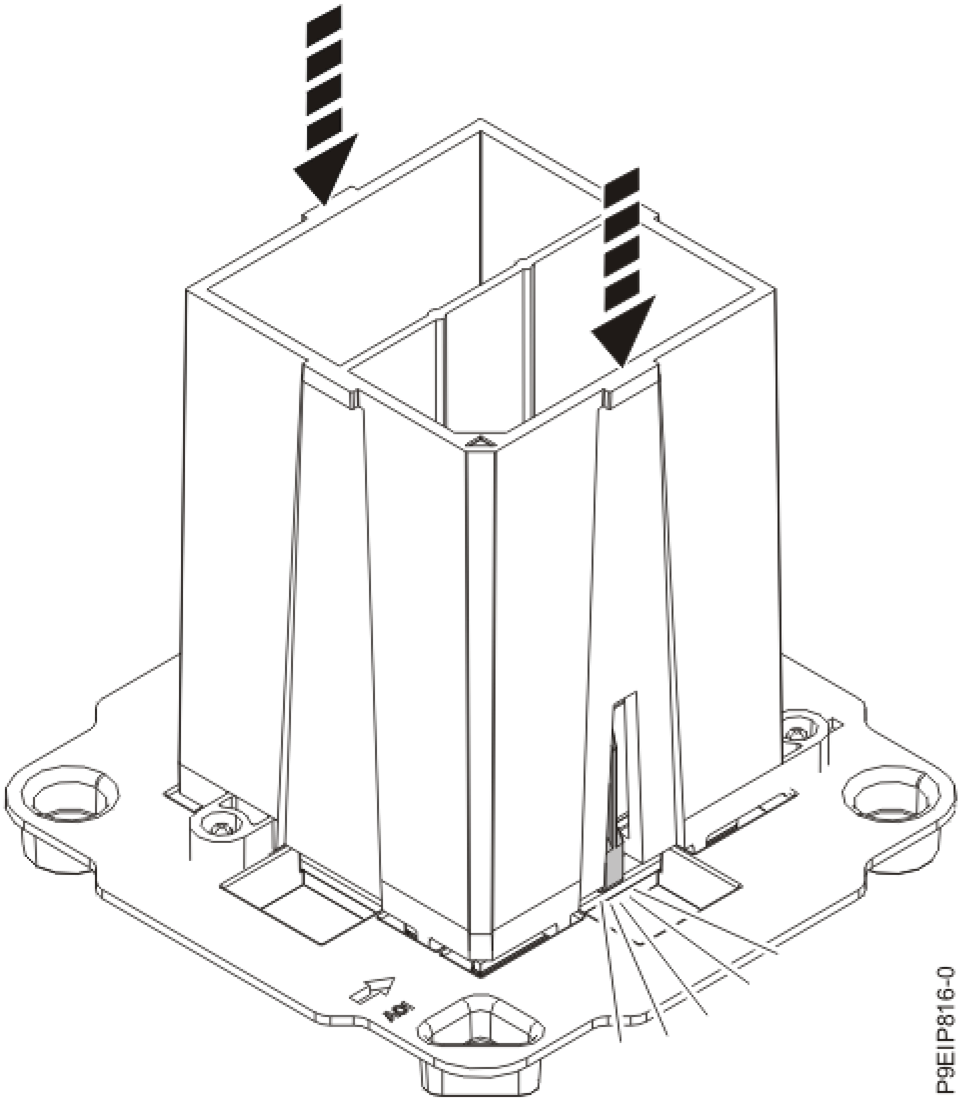

Figure 100. Locking the system processor module into the tool

5. Lift the system processor module from the packaging tray as shown in the following figure.

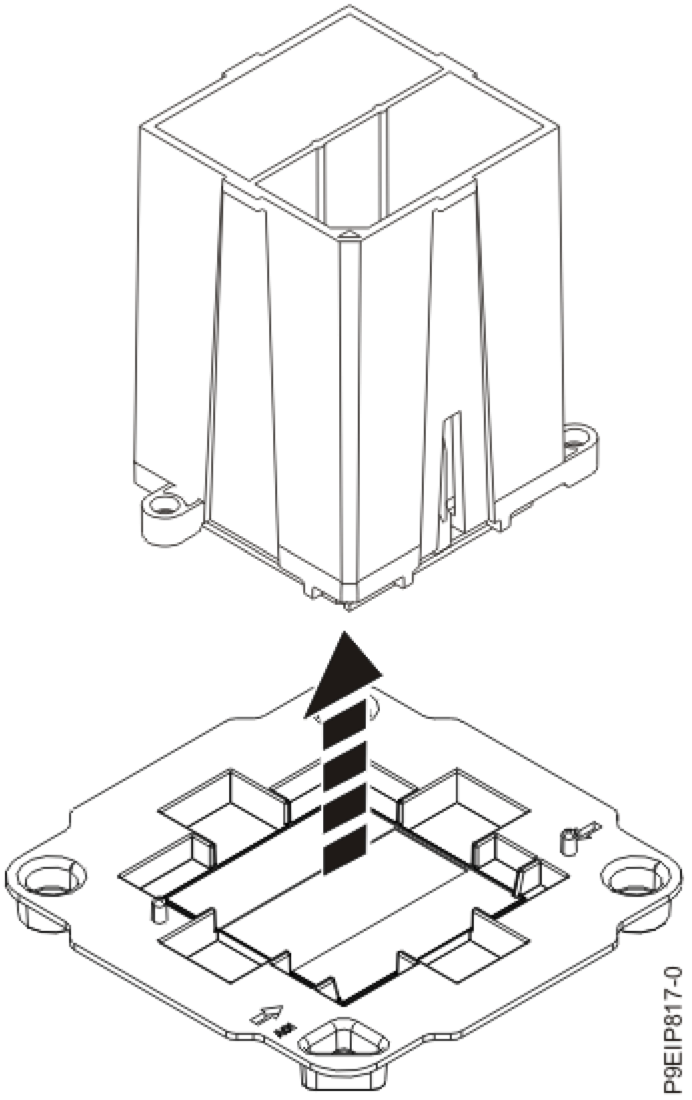

Figure 101. Lifting the system processor module from the packaging tray

6. Lower the tool and system processor module onto the socket. Align the beveled corner **(A)** of the tool with the beveled corner on the socket as shown in the following figure.

   Ensure that the two guide pins **(C)** are inserted into the alignment holes **(B)** on each side of the tool. Use care to lower the tool evenly without tilting the tool. Do not attempt to slide the tool and the system processor module in any direction while the system processor module is touching the socket. If the tool and the system processor module are not aligned with the guide pins, lift the tool and the system processor module and reposition them.

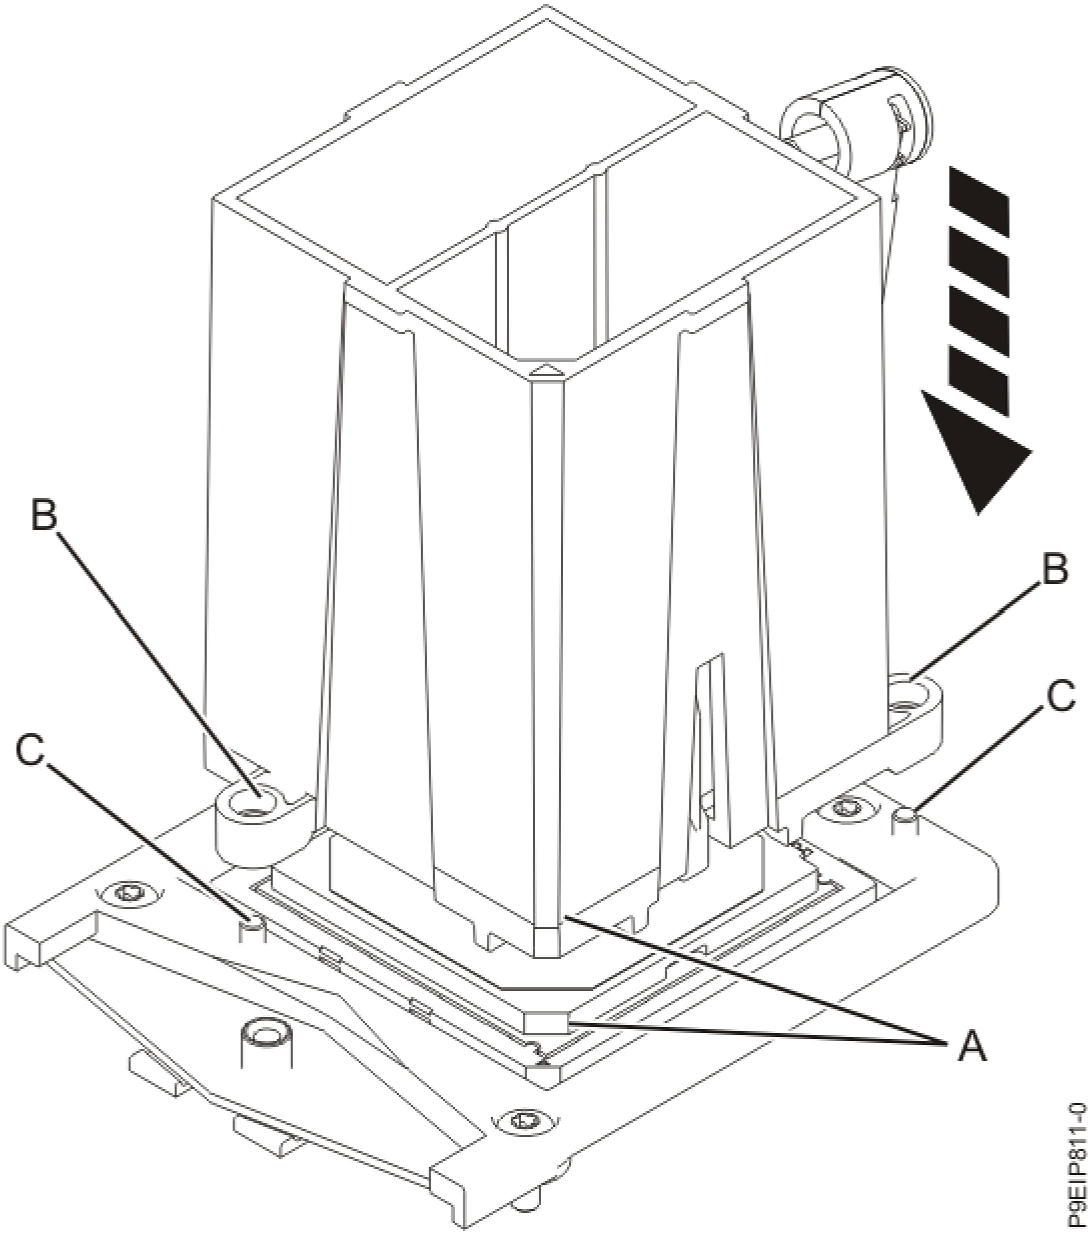

Figure 102. Installing the system processor module

7. After the tool and system processor module holes and guide pins are properly aligned, squeeze and hold the two blue release tabs **(A)** together until a firm stop is reached as shown in the following figure.

   Then, lift the tool off the system processor module.

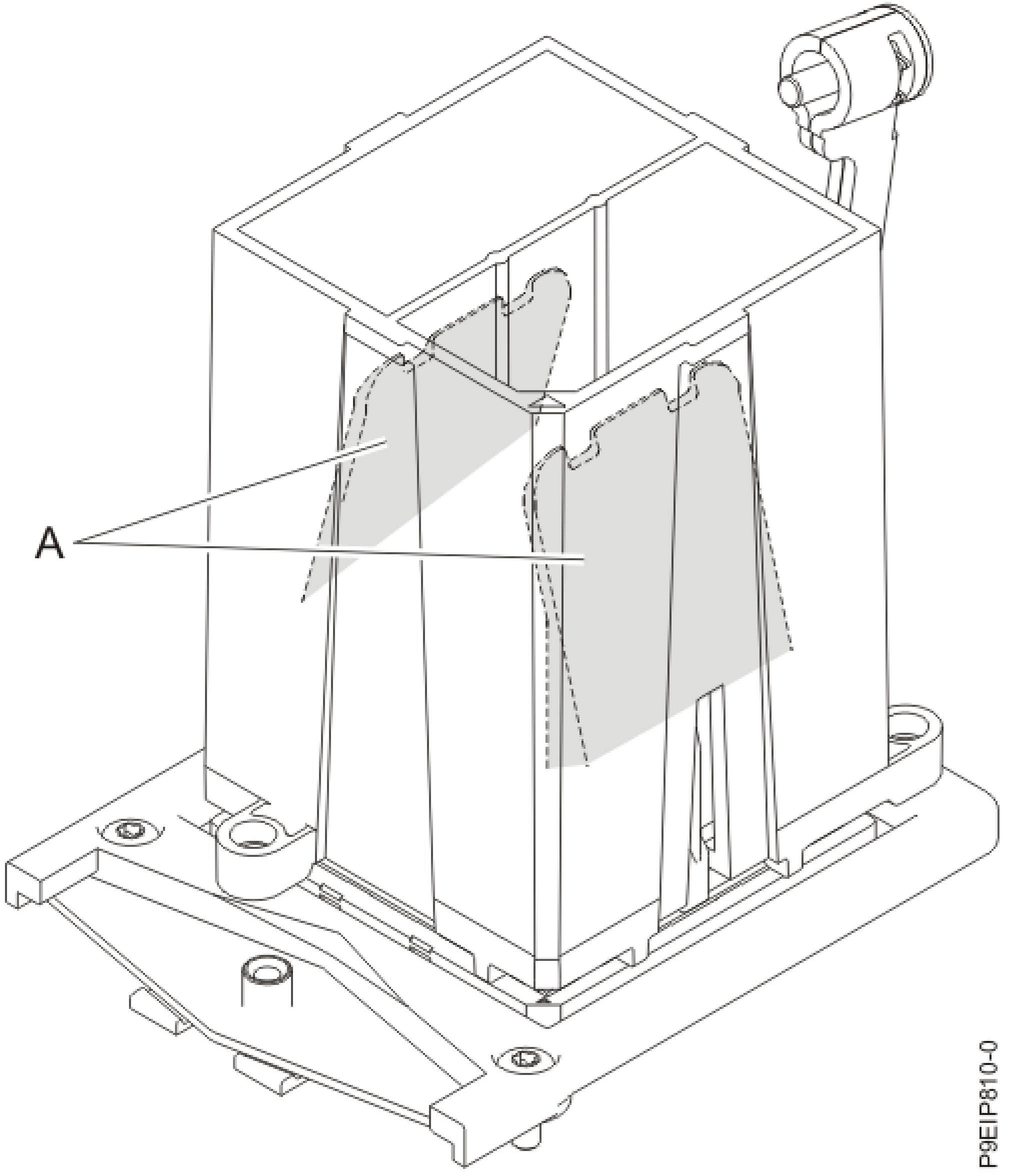

Figure 103. Removing the system processor module tool

8. Inspect the thermal interface material (TIM) for visible signs of damage. If you see folds, tears, bends, or if you have doubts about the TIM, replace it.

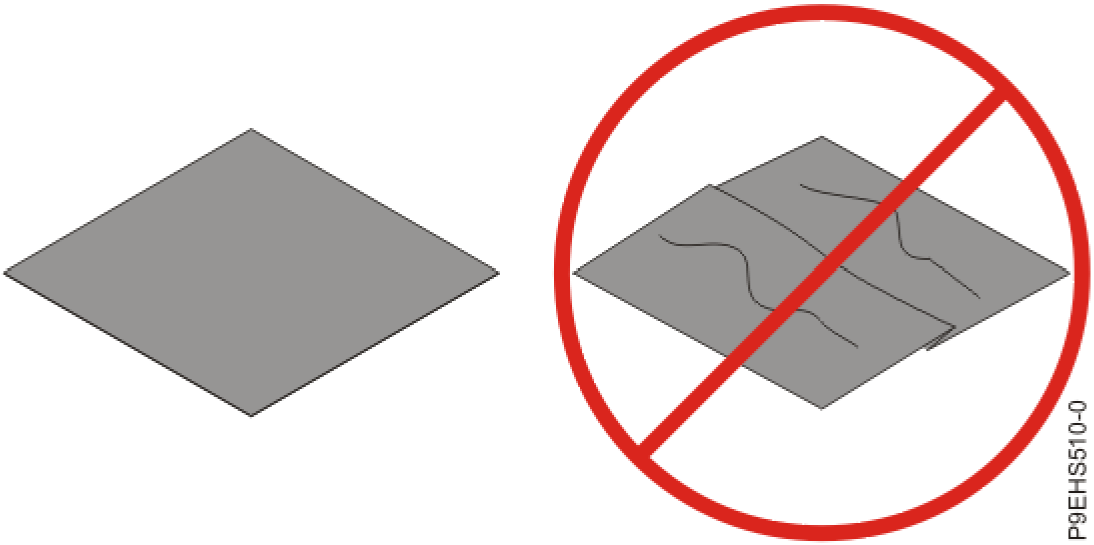

Figure 104. Inspecting the thermal interface material

9. Choose one of the following repair options:

Option	Description
Is the TIM damaged?	It is damaged. Proceed to step "10" on page 82 to replace the TIM and install the existing heat sink.
Is the TIM OK?	It is not damaged and can be reused. Proceed to step "12" on page 82 to reuse the TIM and install the existing heat sink.

10. Use this step to install a new TIM and reuse the existing heat sink.
    a) Open the TIM packaging and carefully remove the TIM, holding it by the edges of the carrier strip and holding it away from the shipping container.
    b) Remove the protective film from the clear carrier strip by using the supplied tweezers.

       Note: The TIM must remain flat. Small wrinkles are acceptable, but folds are not acceptable.
    c) Using the tweezers, remove the TIM from the carrier strip and center it onto the system processor module.

       The TIM has no preferred up side. The TIM can be placed on the system processor module and centered..
11. Continue with step "13" on page 82.
12. Use this step to reuse the existing undamaged TIM and heat sink.
    a) Using the tweezers, move the old TIM from the clean, dry surface and center it onto the new system processor module.

       The TIM has no preferred up side. The TIM can be placed on the system processor module and centered.
13. Carefully lower the heat sink over the system processor module, ensuring that the holes in the heat sink align with the two guide pins **(A)** on the socket, as shown in the following figure.

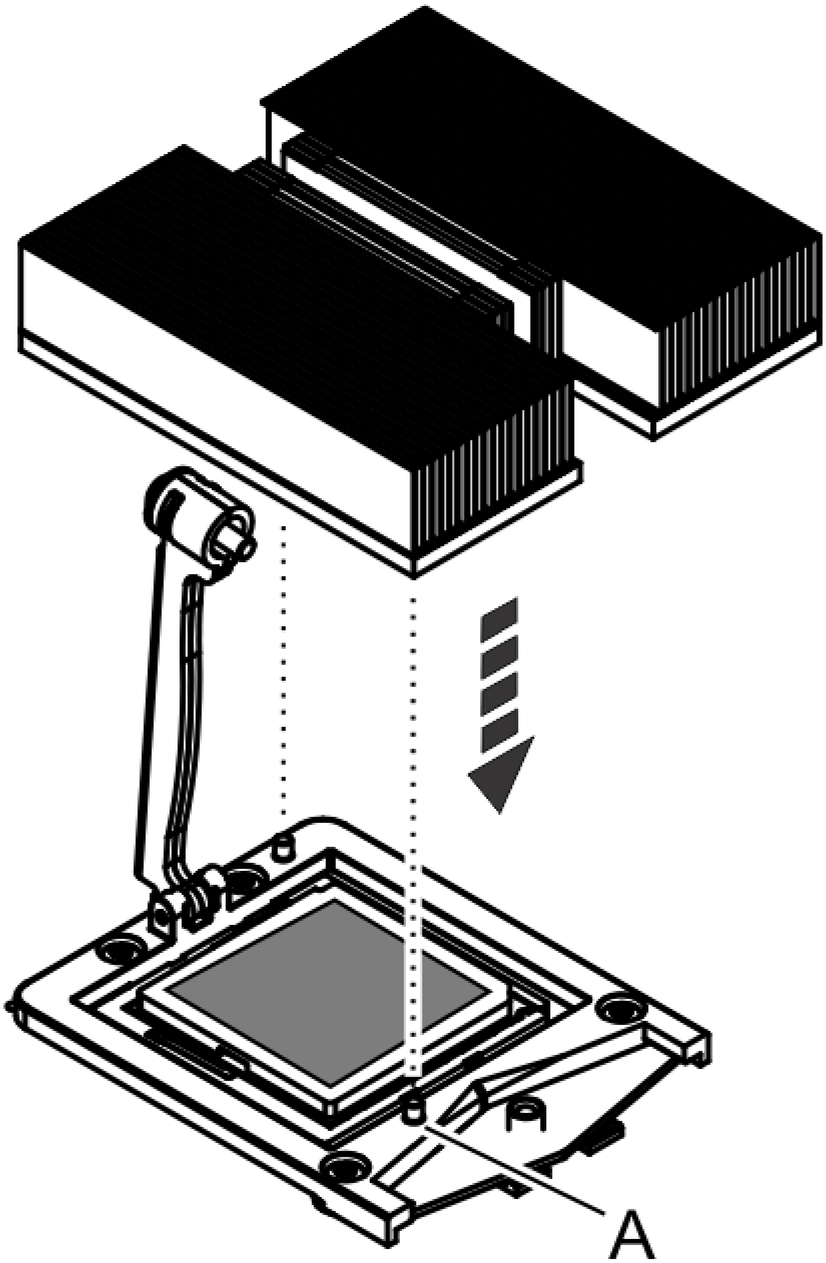

Figure 105. Installing the heat sink

14. Move the load arm **(A)** into position over the heat sink **(B)** and tighten the load arm screw with a T20 hexalobular driver, as shown in the following figure.

    Note: Do not over tighten the load arm screw.

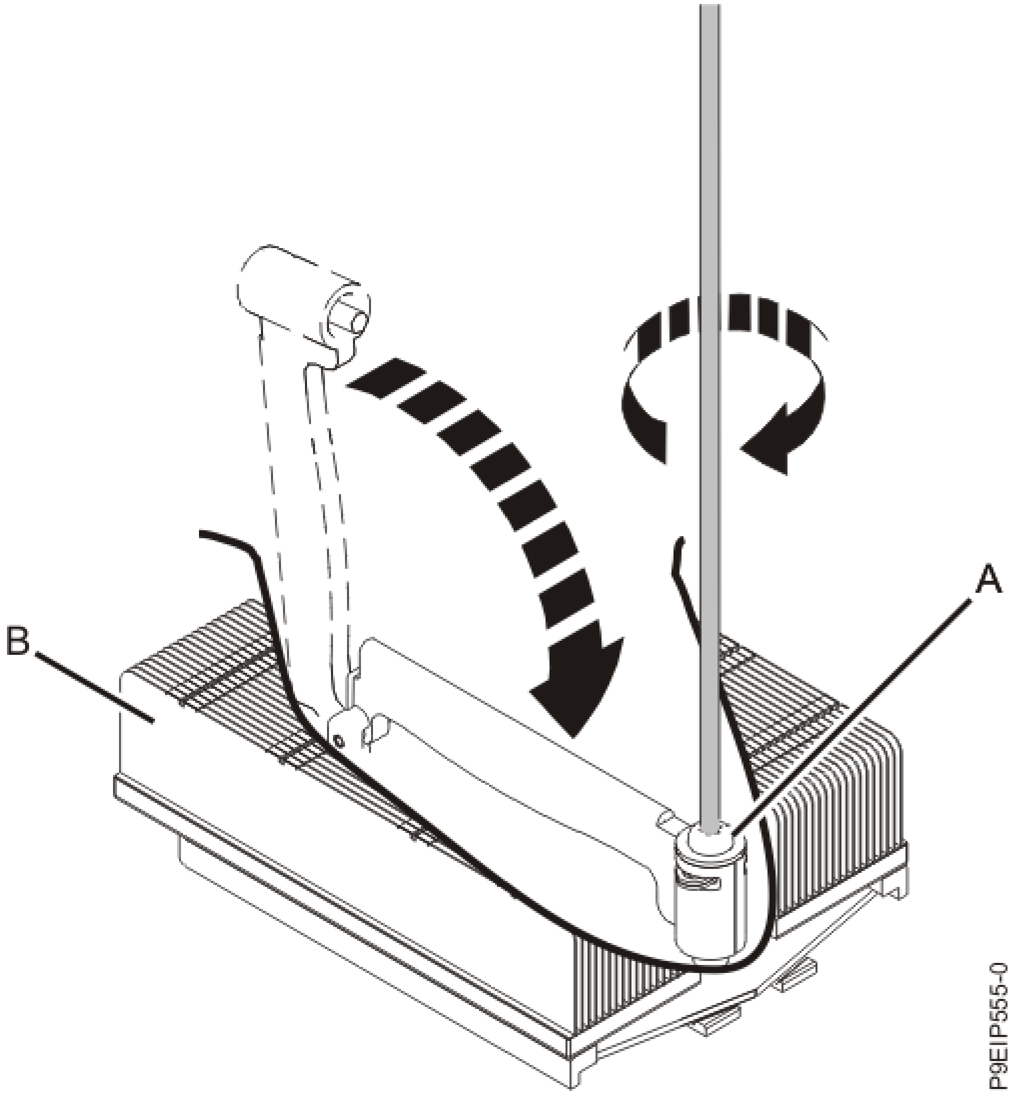

Figure 106. Tightening the load arm screw

15. Lightly grip the system processor module that you replaced by the edges and lift it off the shipping cover. Align the beveled corner of the module **(A)** to the corner of the tray with the triangle **(B)** and place it in the tray, as shown in the following figure.

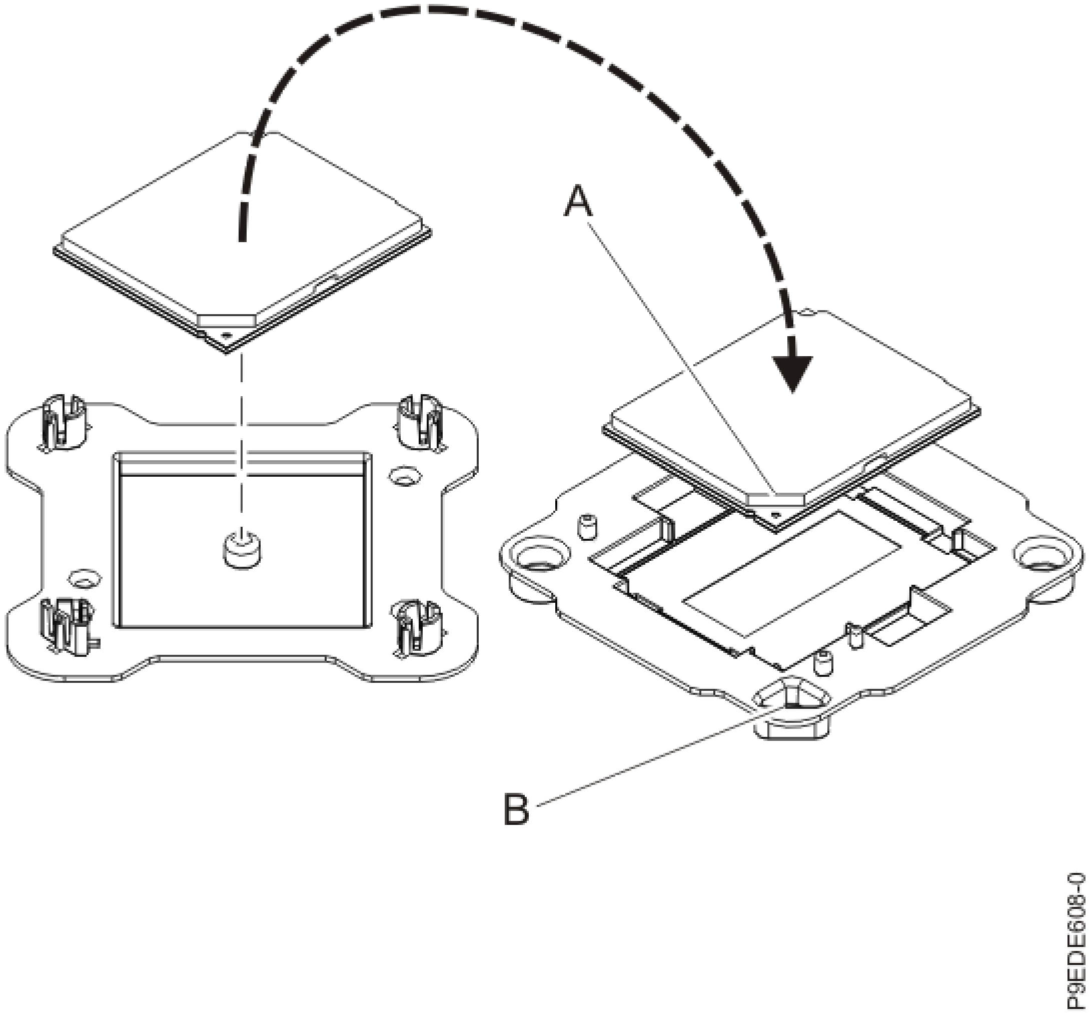

Figure 107. Placing the system processor module into the shipping tray

16. Replace the system backplane into the rear of the system.
    a) Ensure that the two system backplane levers are open.
    b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

    Important:

    - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
    - Ensure that the system backplane is fully seated and is all the way into the system.
    - You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

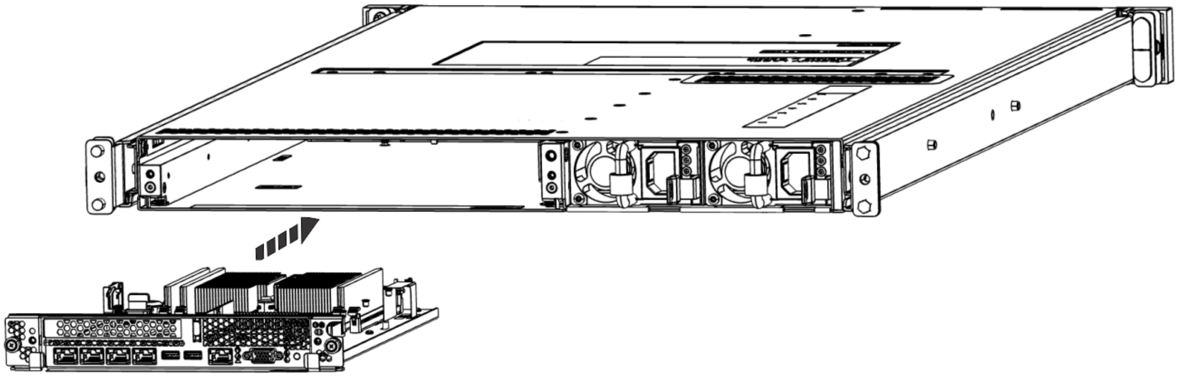

Figure 108. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

17. Power on the system for operation.

For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the time-of-day battery in the 7063-CR2

To remove and replace the time-of-day battery in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Power off the system.

   For instructions, see "Stopping the 7063-CR2 system" on page 99.

2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

   - Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
   - When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
   - If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

3. Remove the system backplane from the rear of the system.

   a) Label and remove the two power cables.

   For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.

   b) Label and remove the signal cables from the rear of the system.

   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

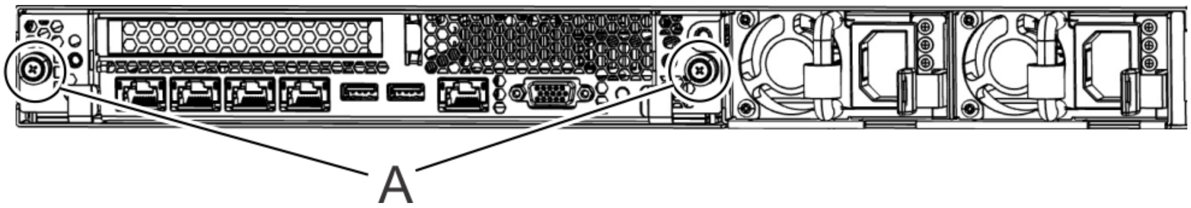

Figure 109. Removing the system backplane screws

d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

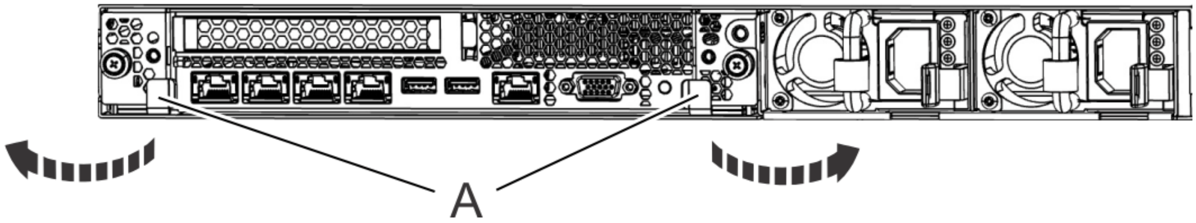

Figure 110. Unlatching the system backplane

e) Support the system backplane by the bottom as you slide it from the system.

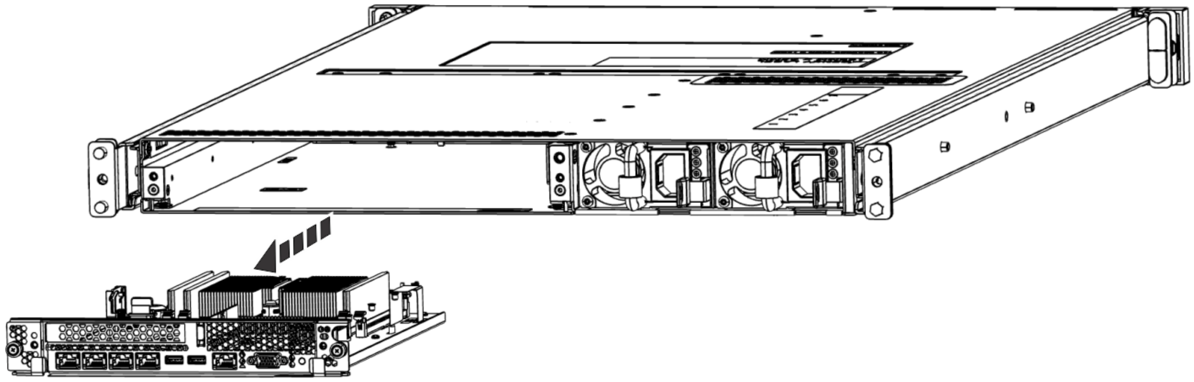

Figure 111. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

4. Remove the time-of-day battery **(A)** by using your thumb to press the spring latch toward the back of the system to release the time-of-day battery. Lift the time-of-day battery from the battery socket.

   When you remove the time-of-day battery, do not use a metallic tool to disengage it from its slot.

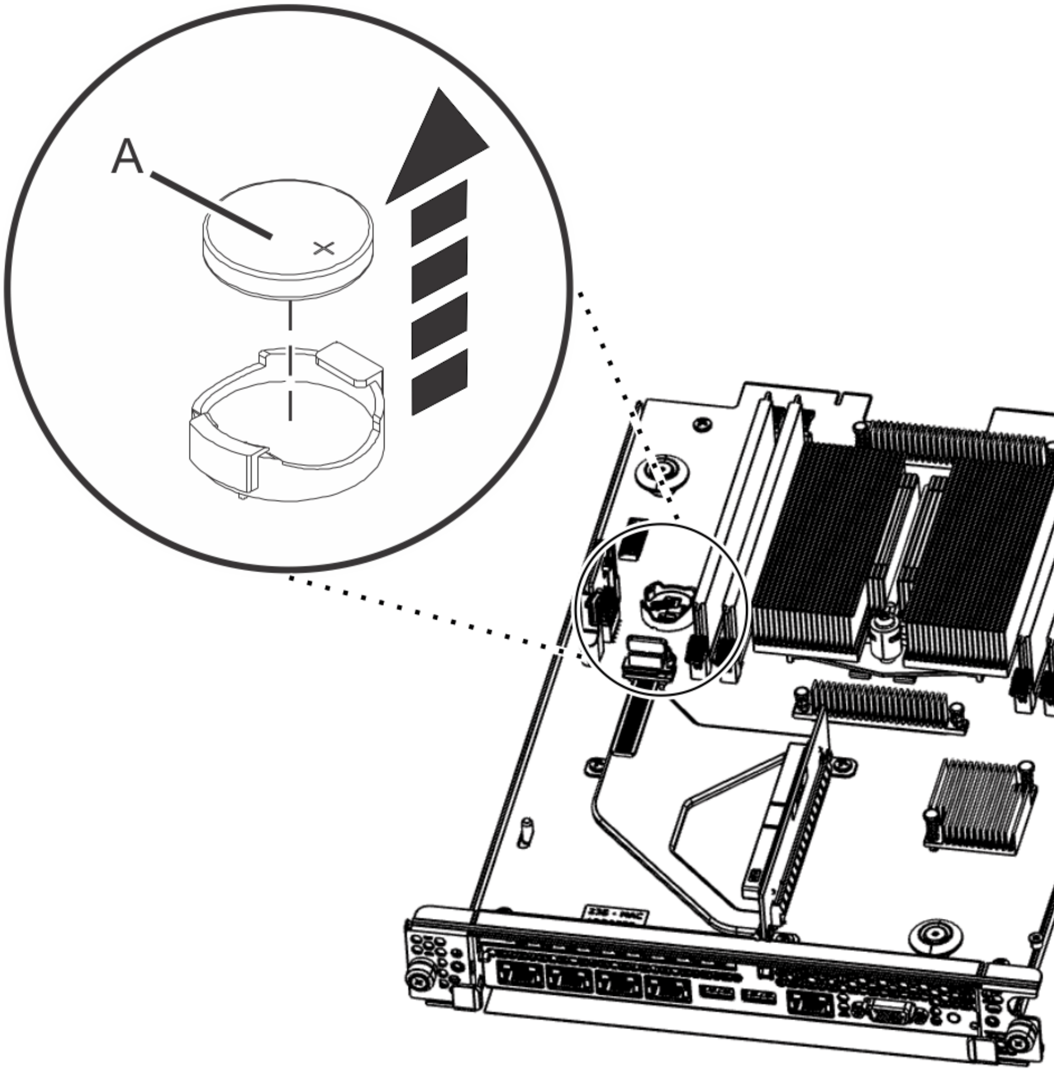

Figure 112. Removing the time-of-day battery

5. To replace the time-of-day battery, use your thumb to press the battery socket spring latch and replace the time-of-day battery.

   The orientation of the + of the battery is up.

6. Replace the system backplane into the rear of the system.

   a) Ensure that the two system backplane levers are open.

   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

      Important:

      - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
      - Ensure that the system backplane is fully seated and is all the way into the system.
      - You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

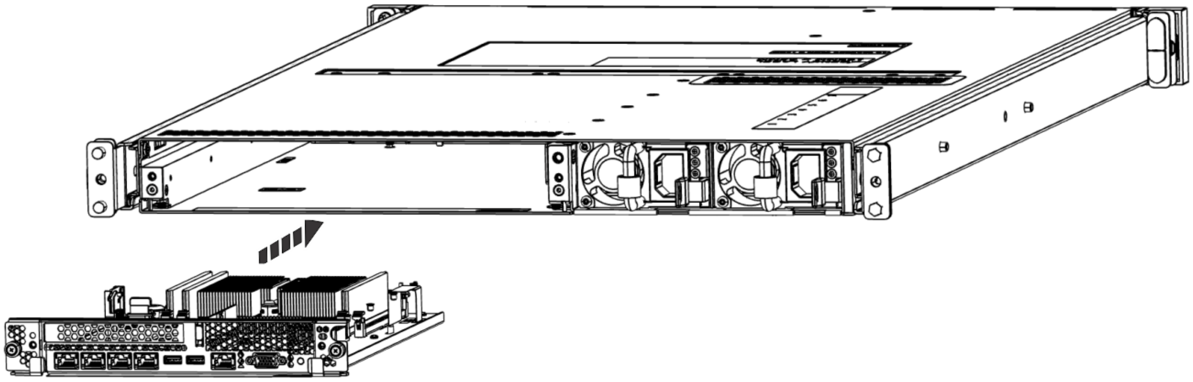

Figure 113. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

7. The date and time needs to be set before the operating system starts to avoid issues.

   Follow these petitboot steps:

   a) From the petitboot menu, select **Exit to shell**.

   b) Run the following two commands to check the date and time:

```
date
ipmitool sel time get
```

   c) To set the correct date and time in UTC use the following format:

```
date -s YYYY.MM.DD-HH:MM
```

   d) Run the following ipmitool command to set the system entry log (SEL) date and time to the new values:

```
ipmitool sel time set now
```

   e) Verify that the ipmitool and date commands now show the correct values by running the following commands again:

```
date
ipmitool sel time get
```

   f) To leave the petitboot shell, type **exit**.

8. Power on the system for operation.

   For instructions, see "Starting the 7063-CR2 system " on page 98.

Removing and replacing the trusted platform module in the 7063-CR2

To remove and replace the trusted platform module in the IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Attach the electrostatic discharge (ESD) wrist strap.

The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

Attention:

- Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
- When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
- If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

2. Remove the system backplane from the rear of the system.
   a) Label and remove the two power cables.

      For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.
   b) Label and remove the signal cables from the rear of the system.
   c) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

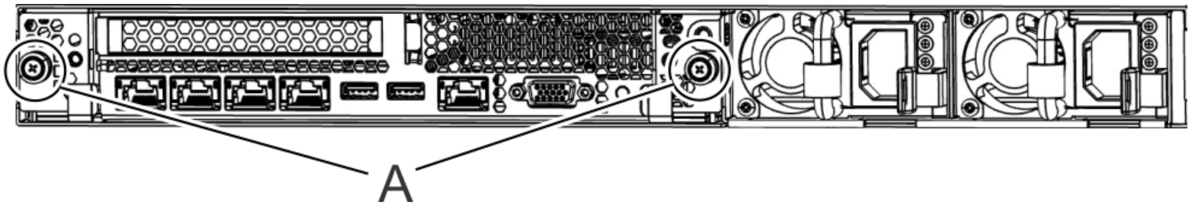

Figure 114. Removing the system backplane screws

   d) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

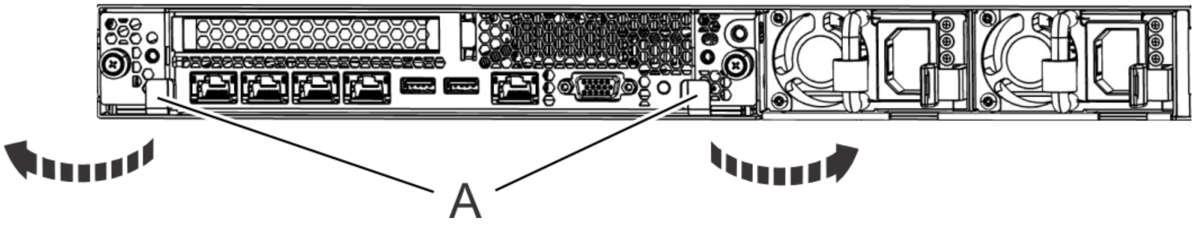

Figure 115. Unlatching the system backplane

   e) Support the system backplane by the bottom as you slide it from the system.

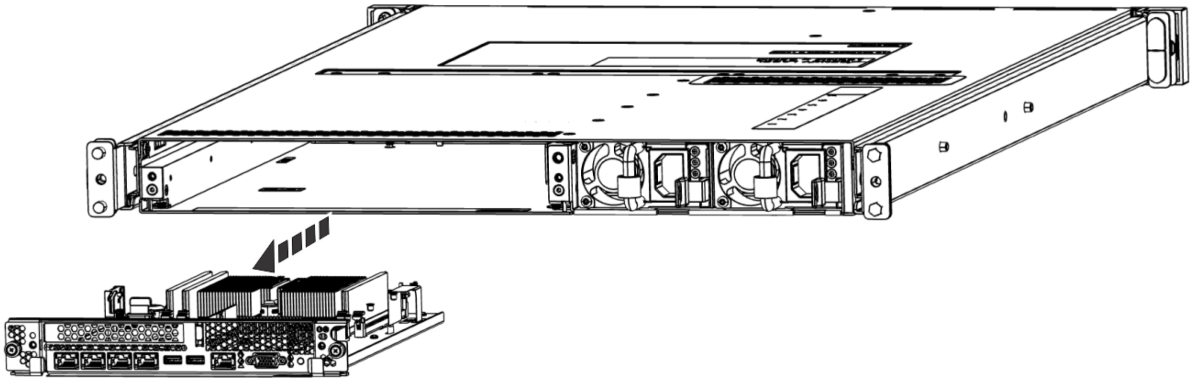

Figure 116. Removing the system backplane

f) Place the system backplane on an ESD surface.

Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

3. Pull small the lever **(A)** slightly away from the trusted platform module to release the module and lift the module straight up from its slot on the system backplane.

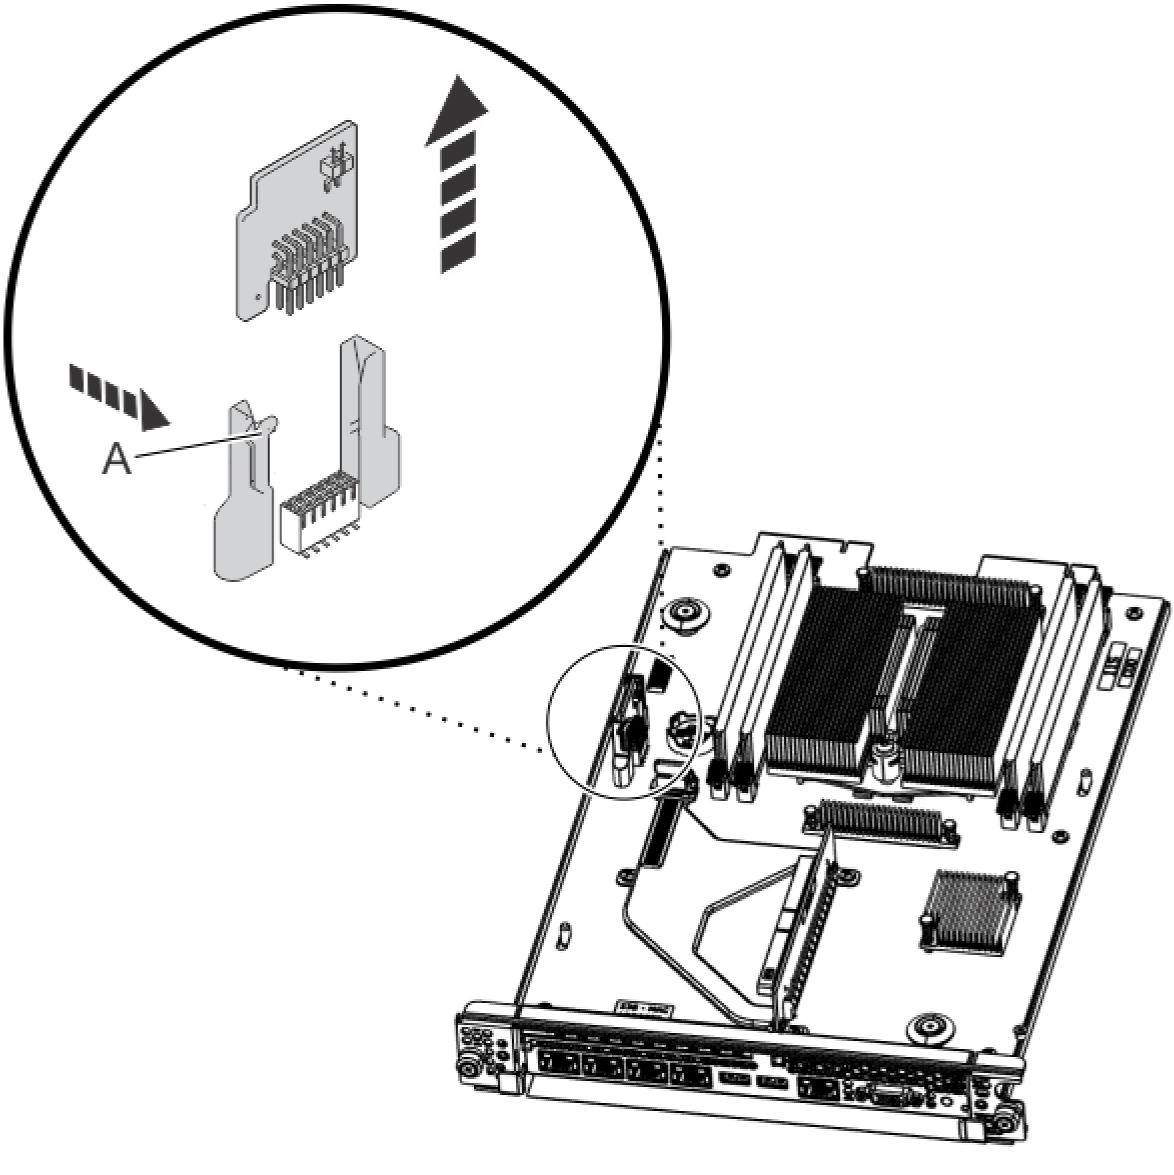

Figure 117. Removing the trusted platform module

4. Align the trusted platform module with the plastic guides and push the trusted platform module straight into the system backplane until it is fully seated and the lever **(A)** clicks into place.
5. Replace the system backplane into the rear of the system.

   a) Ensure that the two system backplane levers are open.

   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

   Important:

   - Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.

- Ensure that the system backplane is fully seated and is all the way into the system.
- You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

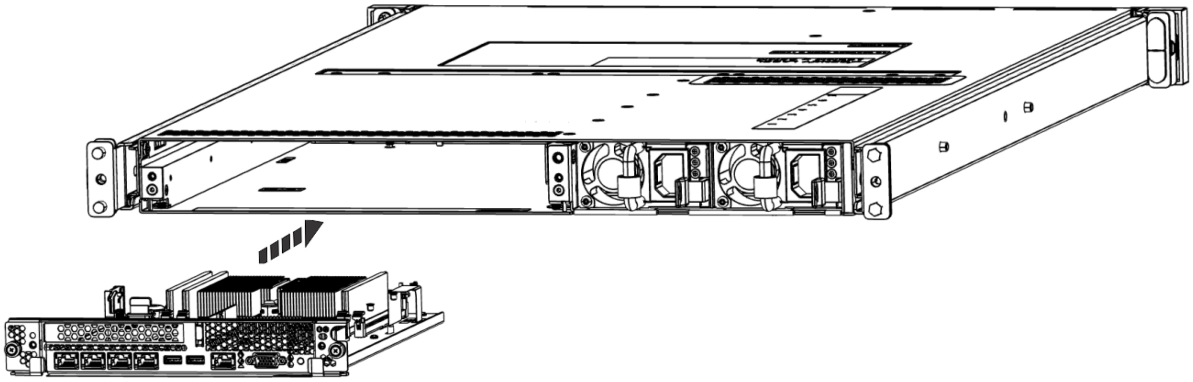

Figure 118. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

f) Using your labels, replace the two power cords at the rear of the system.

For instructions, see “Connecting the power cords to a 7063-CR2 system” on page 101.

6. If secure boot option is enabled before the replacement of the trusted platform module, then the user needs to re-enable the secure boot option by following the steps in Enabling secure boot on 7063-CR2 HMC.

Common procedures for servicing the 7063-CR2

Learn about the common procedures related to removing and replacing parts in the IBM Power Systems HMC (7063-CR2) system.

Before you begin

Observe these precautions when you are installing, removing, or replacing features and parts.

About this task

These precautions are intended to create a safe environment to service your system and do not provide steps for servicing your system. The installation, removal, and replacement procedures provide the step-by-step processes that are required to service your system.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer’s DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer’s DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

  To Disconnect: 1) Turn off everything (unless instructed otherwise). 2) For AC power, remove the power cords from the outlets. 3) For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source. 4) Remove the signal cables from the connectors. 5) Remove all cables from the devices.

  To Connect: 1) Turn off everything (unless instructed otherwise). 2) Attach all cables to the devices. 3) Attach the signal cables to the connectors. 4) For AC power, attach the power cords to the outlets. 5)

For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP. 6) Turn on the devices.

- Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment–personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet if provided, unless the earthquake option is to be installed.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).
- Stability hazard:
  - The rack may tip over causing serious personal injury.
  - Before extending the rack to the installation position, read the installation instructions.
  - Do not put any load on the slide-rail mounted equipment mounted in the installation position.
  - Do not leave the slide-rail mounted equipment in the installation position.
- Each rack cabinet might have more than one power cord.
  - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
  - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer's DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (R001 part 1 of 2)

(R001 part 2 of 2):

CAUTION:

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection.

To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.

- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack or if the rack is not bolted to the floor. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.

- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

Procedure

1. If you are installing a new feature, ensure that you have the software that is required to support the new feature. See IBM Prerequisite.
2. If you are installing or replacing something that might put your data at risk, ensure, wherever possible, that you have a current backup of your system or logical partition (including operating systems, licensed programs, and data).
3. Review the installation or replacement procedure for the feature or part.
4. Note the significance of color on your system.

   Blue on a part of the hardware indicates a touch point where you can grip the hardware to remove it from or install it in the system, or open or close a latch.
5. Ensure that you have access to a medium flat-blade screwdriver, a Phillips screwdriver, and a pair of scissors.
6. If parts are incorrect, missing, or visibly damaged, do the following steps:
   - If you are replacing a part, contact the provider of your parts or next level of support.
   - If you are installing a feature, contact one of the following service organizations:
     - The provider of your parts or next level of support.
     - In the United States, the IBM Rochester Manufacturing Automated Information Line (R-MAIL) at 1-800-300-8751.

   In countries and regions outside of the United States, use the following website to locate your service and support telephone numbers:

   ```
   http://www.ibm.com/planetwide
   ```
7. If you encounter difficulties during the installation, contact your service provider, your IBM reseller, or your next level of support.
8. For thermal performance, ensure that the top cover is on when the system is running.

9. If you are installing new hardware in a logical partition, you need to understand and plan for the implications of partitioning your system. For information, see Logical Partitioning.

Identifying the 7063-CR2 system that contains the part to replace

Learn how to determine which system has the part you want to replace.

LEDs on the 7063-CR2 system

Use this information as a guide to the LEDs on the IBM Power Systems HMC (7063-CR2) system.

The LEDs indicate various system statuses. If the part does not have a problem indicator LED, you can use a troubleshooting program such as **`impitool`** to identify the issue.

The front control panel LEDs are shown in the following figure.

- The green LED **(6)** indicates the power status (on or off). The LED flashes when the BMC is at standby. The LED is solid when the system is running.
- The blue identify LED **(8)** identifies the system to be serviced.
- The amber LED **(7)** indicates a system fault.
- The fan LEDs **(1)** - **(5)** indicate an issue with the corresponding fan.

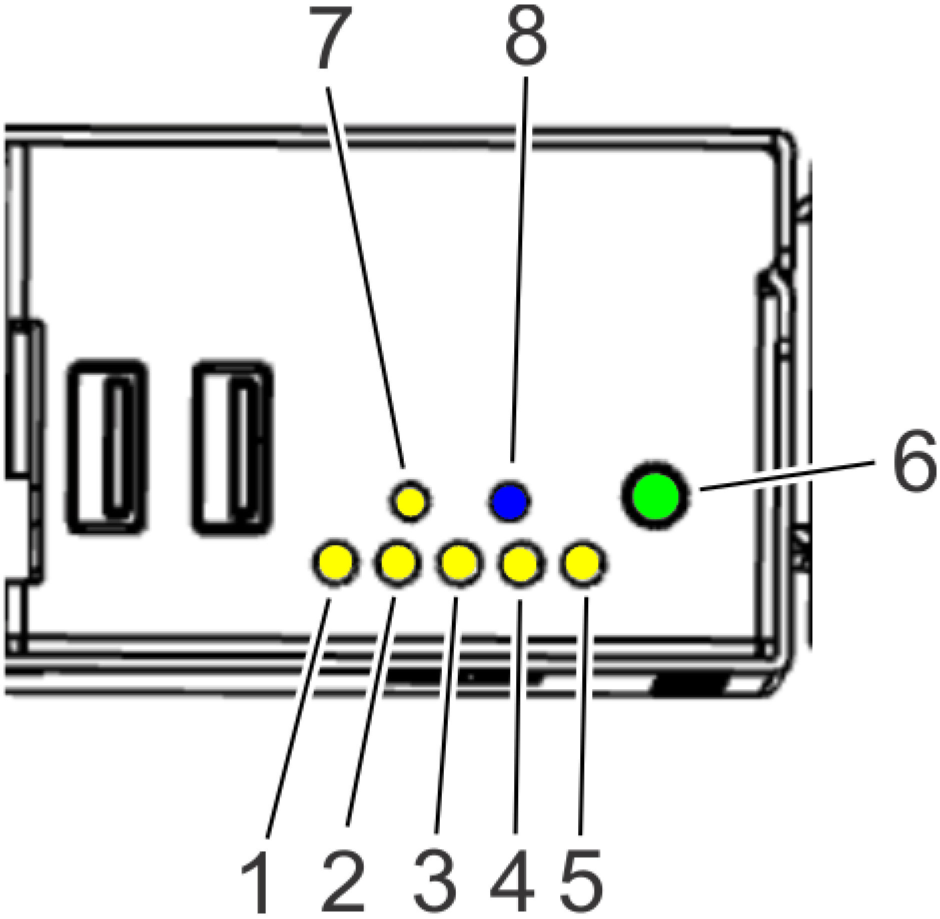

Figure 119. Control panel LEDs

The drive LEDs are shown in the following figure.

- The green LED indicates the power status (on or off).
- The amber LED flashes when there is activity.

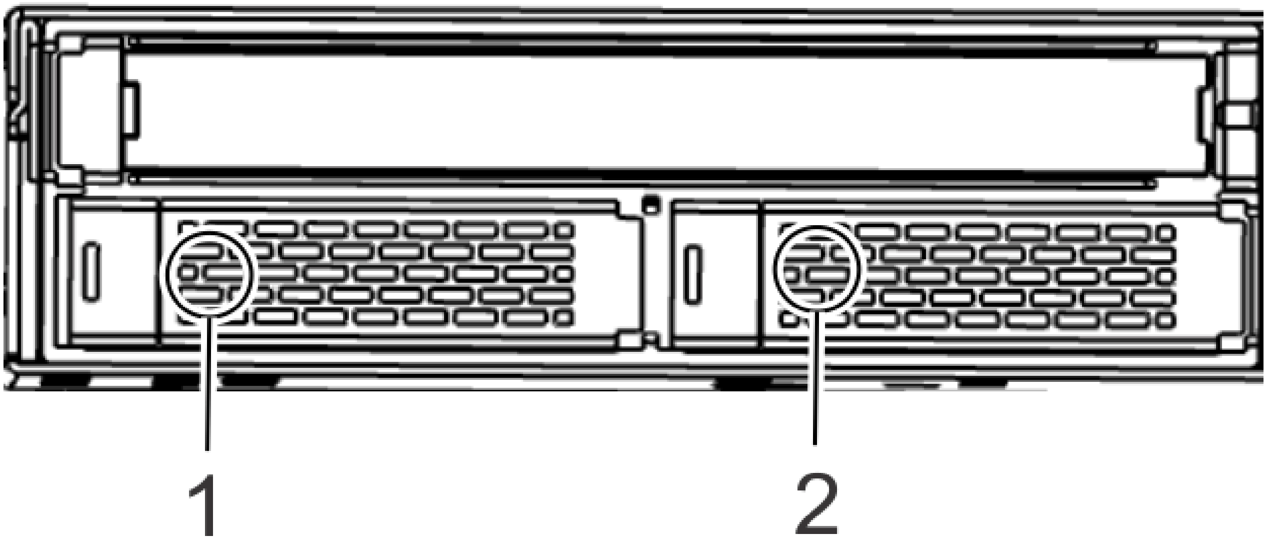

Figure 120. Drive LEDs

LEDs are also on the rear of the system; see the following figure.

- The blue identify LED **(1)** identifies the system to be serviced.
- The amber LED **(2)** indicates a system fault.

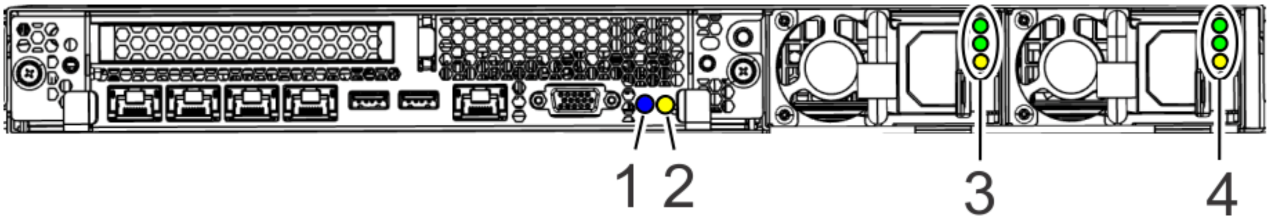

Figure 121. LEDs on the rear of the system

Power supply LEDs **(3)** and **(4)** can indicate the following states:

- The top green LED indicates AC power (on or off). The LED is solid during system standby.
- The middle green LED indicates DC power (on or off). The LED flashes during system standby.
- The bottom amber LED indicates a power fault.

Identifying the 7063-CR2 that needs servicing

Use the Intelligent Platform Management Interface (IPMI) program to turn on the blue identify LED to help you find the IBM Power Systems HMC (7063-CR2) system that needs servicing.

Procedure

You can use the following command to activate the blue system identify LED:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis identify on
```

To turn off the blue system identify LED, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis identify
off
```

To check the status of the blue system identify LED, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> chassis identify
status
```

Also visually check the LED.

Starting and stopping the 7063-CR2

Learn how to start and stop the IBM Power Systems HMC (7063-CR2) system to perform a service action or system upgrade.

Starting the 7063-CR2 system

You can use the power button to start the IBM Power Systems HMC (7063-CR2) systems.

About this task

⚠ **Attention:** For safety, airflow purposes and thermal performance, the service access cover must be installed and fully seated before you power on the system.

You can use this procedure to power on the system, or you can use a console and the IPMI tool to power on the system.

Procedure

1. Before you press the power button, ensure that the power supplies are connected to the system unit and that the power cables are connected to a power source.
2. Press the power button **(6)** shown in the following figure.

   The power-on light stops flashing and remains on, indicating that the system power is on.

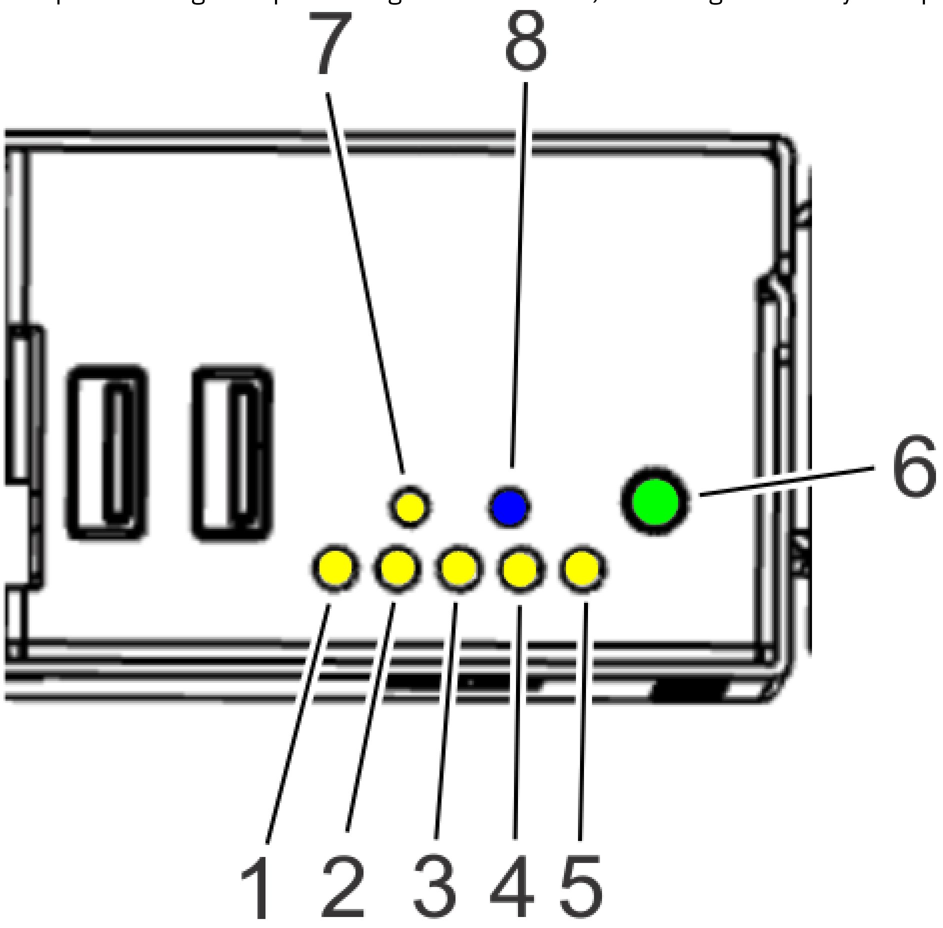

Figure 122. Power switch for the 7063-CR2 system

What to do next

If you press the power button and the system does not start, contact your next level of support or your service provider.

Stopping the 7063-CR2 system

To stop the IBM Power Systems HMC (7063-CR2) systems, complete the steps in this procedure.

Procedure

You can use the **`hmcshutdown`** command to stop and power down the system.

For example, the following command shuts down the system now.

```
hmcshutdown -t now
```

Drive commands for the 7063-CR2 system

Learn about the drive commands for the IBM Power Systems HMC (7063-CR2) system.

The system uses the `arcconf` command. The command is included in Petitboot. In Petitboot shell, run the command including the `bin` folder; for example:

```
/bin/arcconf
```

When you run the `arcconf` command from the operating system, you must be logged in as root.

To list the drive and device configurations, use the following `arcconf` command.

```
arcconf getconfig 1 pd
```

Note the Channel and Device numbers of the desired drive. In the following example, the channel number is 0 and the device number is 1.

```
Reported Channel,Device(T:L)       : 0,1(1:0)
```

Identify the fault LED on the drive using the following command. The command uses the Channel and Device numbers from the **`arcconf getconfig`** command:

```
arcconf identify 1 device 0 1
```

Take the drive off line. Run the following **`arcconf setstate`** command. The command uses the Channel and Device numbers from the **`arcconf getconfig`** command:

```
arcconf setstate 1 device 0 1 ddd
```

Creating the virtual drive on the 7063-CR2 system

Follow the steps in this procedure in the unlikely event that you need to re-create the virtual drive on the IBM Power Systems HMC (7063-CR2).

About this task

This procedure assumes that:

- The system was erroneously shipped without an already created and preinstalled virtual disk.
- The existing virtual disk is somehow damaged and needs to be re-created.

These are rare events.

CAUTION: This procedure results in loss of data. This procedure must be used only if the HMC operating system is either not installed, incorrectly installed, or corrupted.

To create the RAID1 logical drive and to rebuild the content on that drive, run the following commands from the Petitboot shell:

1. Create the RAID1 logical drive:

   ```
   /bin/arcconf create 1 logicaldrive name “HMC Disk” max 1 0,0 0,1
   ```

2. Display the logical drive to confirm that is was created:

   ```
   /bin/arcconf getconfig 1 ld
   ```

3. Enable the automatic rebuild of the logical drive:

   ```
   /bin/arcconf setcontrollerparam 1 spareactivationmode 1
   ```

If you need to remove the RAID1 logical drive:

```
/bin/arcconf delete 1 logicaldrive 0
```

Sensor status

You can check the sensor status to quickly determine the general health of the system without using the event codes.

To view the sensor status, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> fru status
```

To view the sensor status and any corresponding event codes, use the following command:

```
openbmctool -U <username> -P <password> -H <BMC IP address or BMC host name> fru status -v
```

Sensors that have a status of **present** and **functional** do not require a service action. Sensors that have a status of **present** and **not functional** require a service action.

Some occurrences of errors in the system might not appear in the sensor status. After you view the sensor status, look for event codes to determine whether a service action is required.

Removing and replacing power cords on a 7063-CR2 system

Learn how to disconnect and connect the power cords on IBM Power Systems HMC (7063-CR2) systems.

Disconnecting the power cords from a 7063-CR2 system

To disconnect the power cords from an IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Before you begin

Note: This system might be equipped with two or more power supplies. If the removing and replacing procedures require the power to be off, then ensure that all power sources to the system are disconnected.

Procedure

1. Identify the system unit that you are servicing in the rack.

   For instructions, see “Identifying the 7063-CR2 system that contains the part to replace” on page 96.
2. Unfasten the hook-and-loop fasteners from the power cords.

3. Label and disconnect the power cords from the system unit as shown in the following figure.

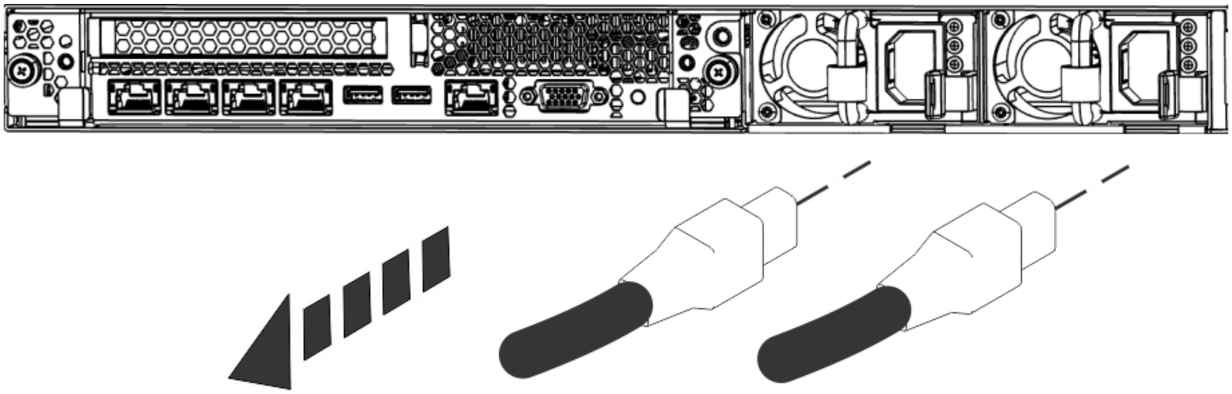

Figure 123. Removing the power cords from the system

Connecting the power cords to a 7063-CR2 system

To connect the power cords to an IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

Procedure

1. Using your labels, reconnect the power cords to the system unit as shown in the following figure.

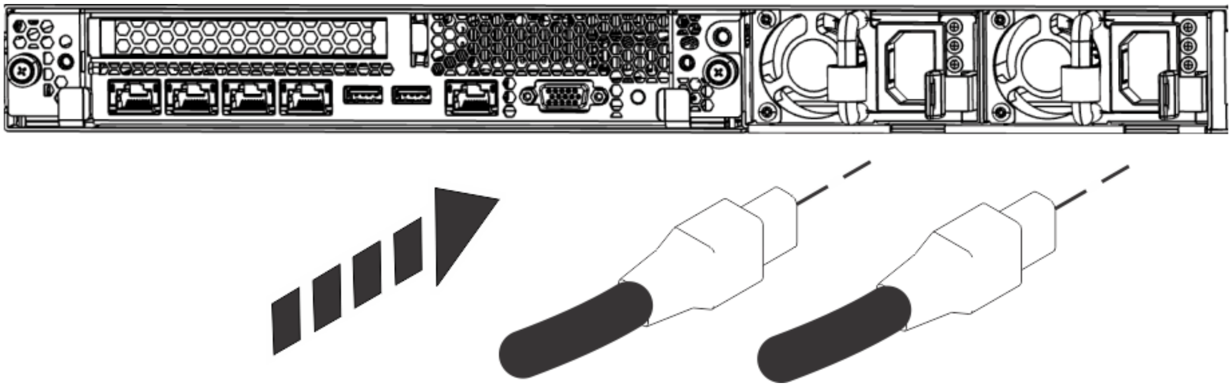

Figure 124. Connecting the power cords to the system

2. Fasten the hook-and-loop fasteners to secure the power cords.

Service and operating positions for the 7063-CR2 system

Learn how to place an IBM Power Systems HMC (7063-CR2) system into the service or operating position.

Placing a 7063-CR2 system into the service position

To place an IBM Power Systems HMC (7063-CR2) system into the service position, complete the steps in this procedure.

Before you begin

The systems must be removed from the rails for servicing some internal parts.

⚠ **CAUTION:** This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

Note: When you move a system out of a rack, ensure that all stability plates are firmly installed to prevent the rack from toppling. Slide only one system out at a time.

Procedure

1. Label and remove the two power cords from the rear of the system.

   For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.
2. Label and remove all cables from the rear of the system.
3. Lighten the system by removing the two power supplies.

   For instructions, see "Removing a power supply from the 7063-CR2 system" on page 51.
4. Lighten the system by removing the system backplane from the rear of the system.

   a) Label and remove the signal cables from the rear of the system.

   b) Loosen the two screws **(A)** on the sides of the system backplane as shown in the following figure.

   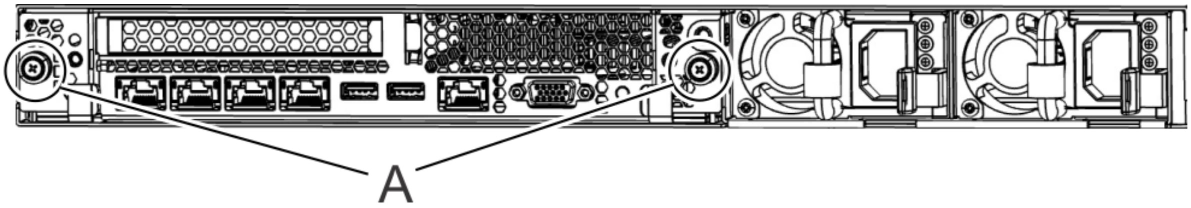

   Figure 125. Removing the system backplane screws

   c) Simultaneously rotate the two levers **(A)** on each side of the system backplane out and to the side to unlock the system backplane from the system.

   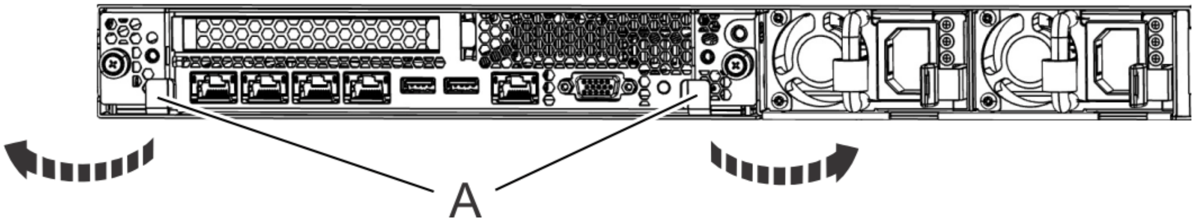

   Figure 126. Unlatching the system backplane

   d) Support the system backplane by the bottom as you slide it from the system.

   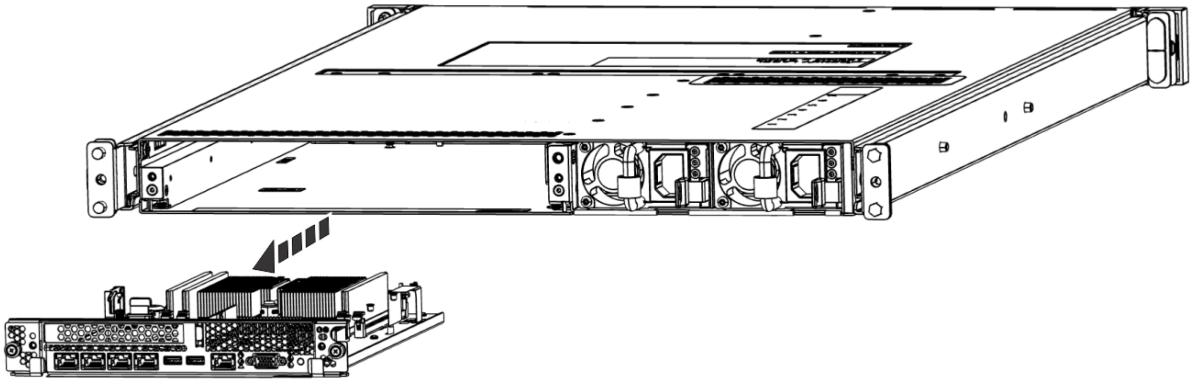

   Figure 127. Removing the system backplane

   e) Place the system backplane on an ESD surface.

   Note: You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.
5. Remove the front screws that secure the system to the rack from both sides of the system as shown in the following figure.

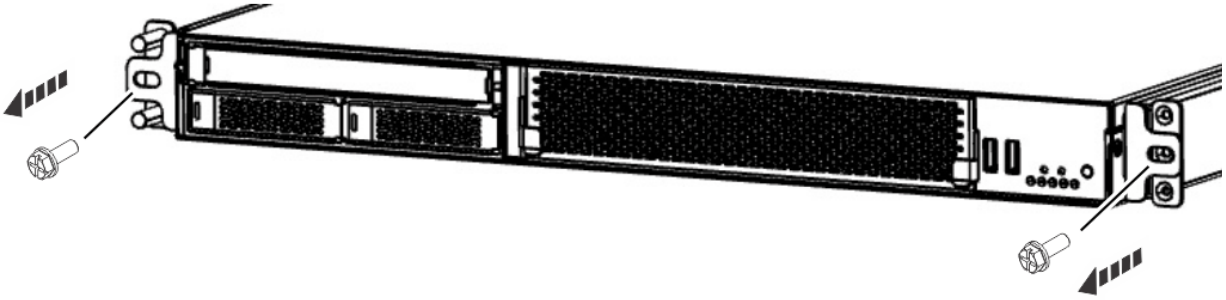

Figure 128. Removing the front screws

6. From the rear of the system, push the system forwards approximately 5 cm (2 in).
7. From the front of the system, while you support the system from underneath, slide the system out of the rack.

   Be careful when removing the system. The rails have no intermediate stopping point. Be sure to support the system from underneath.
8. Carefully set the system on a table with an appropriate ESD surface.

Placing a 7063-CR2 system into the operating position

To place an IBM Power Systems HMC (7063-CR2) system into the operating position, complete the steps in this procedure.

About this task

⚠ **CAUTION:** This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

Procedure

1. Lift the system from the table.
2. Align the rails on each side of the system with the rack slide rails.
3. Push the system into the rack.
4. Fasten the two front screws to secure the system to the rack.

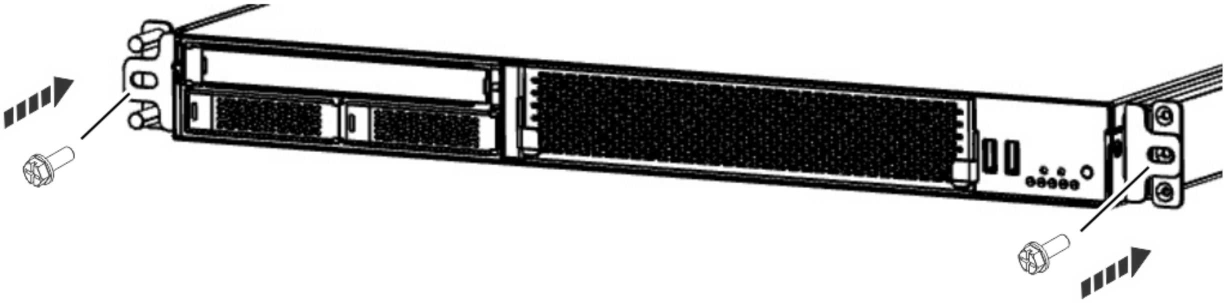

Figure 129. Replacing the front screws

5. If you removed the two power supplies, replace them.

   For instructions, see "Replacing a power supply in the 7063-CR2 system" on page 52.
6. If you removed the system backplane, replace it.
   a) Ensure that the two system backplane levers are open.
   b) Support the system backplane by the bottom as you position the system backplane and insert it into the system until it is fully seated.

      Important:

- Use care when you insert the system backplane so that no damage occurs to the components at the socket edge of the backplane.
- Ensure that the system backplane is fully seated and is all the way into the system.
- You must remove and replace the system backplane at a flat angle. The ventilation holes in the top cover can come into contact with the DIMMs in the system backplane if the insertion of the backplane is at an angle or is rushed. As a result of possible contact, DIMMs can be scratched and can leave residue on the top cover.

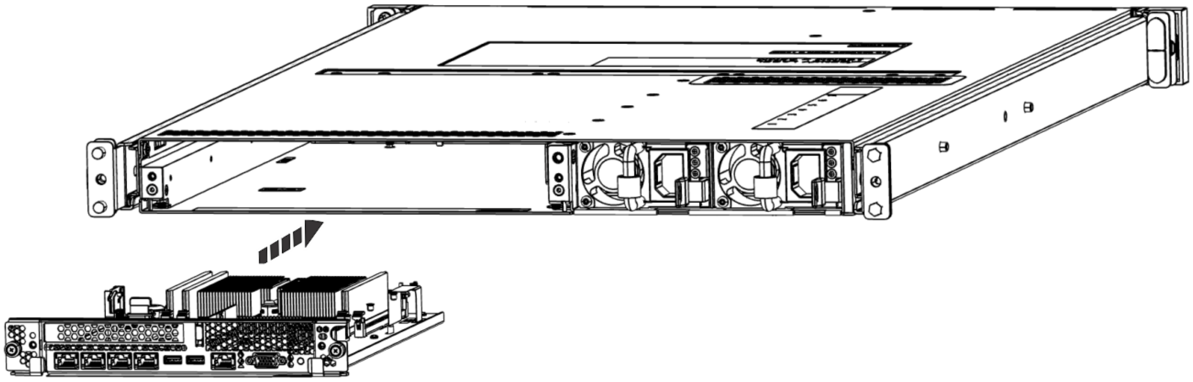

Figure 130. Replacing the system backplane

c) Simultaneously rotate the two levers on each side of the system backplane in to secure the system backplane to the system.

d) Tighten the two screws on the sides of the system backplane.

e) Using your labels, replace the signal cables into the rear of the system.

7. Using your labels, reconnect the cables at the rear of the system unit.
8. Using your labels, replace the two power cords at the rear of the system.

   For instructions, see "Connecting the power cords to a 7063-CR2 system" on page 101.

Removing and replacing covers on a 7063-CR2 system

Learn how to remove and replace the covers on an IBM Power Systems HMC (7063-CR2) system so that you can access the hardware parts or service the system.

Removing the service access cover from a 7063-CR2 system

To remove the service access cover from an IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

⚠ **Attention:** For safety, airflow purposes and thermal performance, the service access cover must be installed and fully seated before you power on the system.

Procedure

1. **If the system is not already powered off and in the service position**, complete these steps: "Placing a 7063-CR2 system into the service position" on page 101
2. Attach the electrostatic discharge (ESD) wrist strap.

   The ESD wrist strap must be connected to an unpainted metal surface until the service procedure is completed, and if applicable, until the service access cover is replaced.

   ⚠ **Attention:**

- Attach an electrostatic discharge (ESD) wrist strap to the front ESD jack, to the rear ESD jack, or to an unpainted metal surface of your hardware to prevent the electrostatic discharge from damaging your hardware.
- When you use an ESD wrist strap, follow all electrical safety procedures. An ESD wrist strap is used for static control. It does not increase or decrease your risk of receiving electric shock when using or working on electrical equipment.
- If you do not have an ESD wrist strap, just prior to removing the product from ESD packaging and installing or replacing hardware, touch an unpainted metal surface of the system for a minimum of 5 seconds. If at any point in this service process you move away from the system, it is important to again discharge yourself by touching an unpainted metal surface for at least 5 seconds before you continue with the service process.

3. Ensure that you removed both power cords from the system. For instructions, see "Disconnecting the power cords from a 7063-CR2 system" on page 100.
4. Remove the rails from both sides of the system.

   a) Remove the two screws that secure the rail to the front of the system.

   b) Slide the rail back and remove the rail from the support pins on side of the system.
5. Remove the 19 screws from the cover as shown in the following figure.

   The system has 4 screws on each side, and 11 screws on the top surface. Use a #2 Phillips screwdriver.

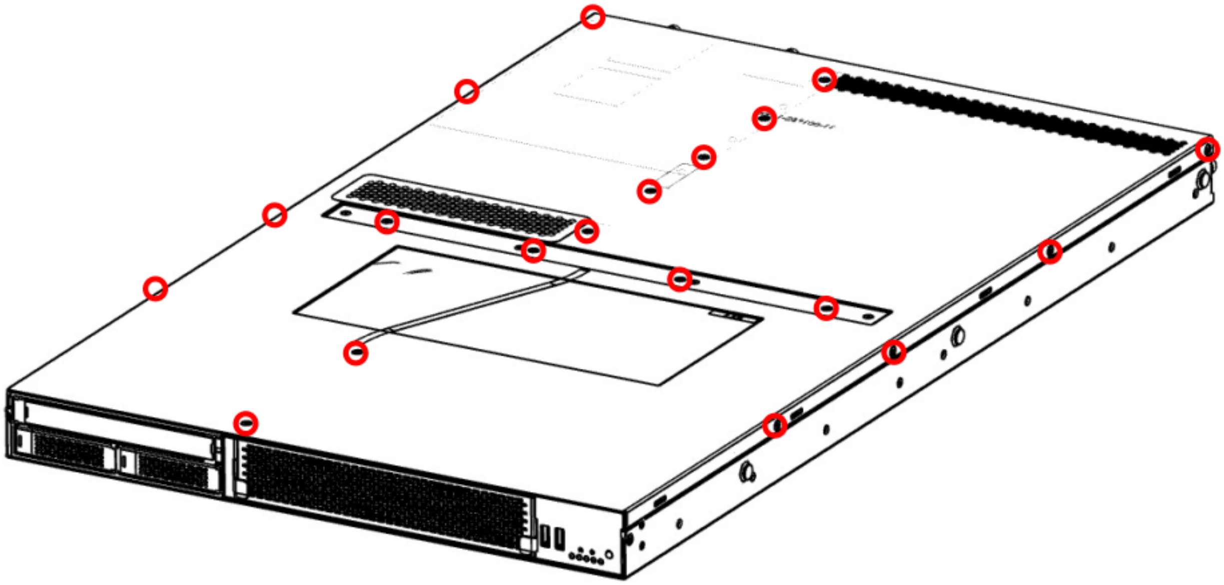

Figure 131. Removing the cover screws

6. Slide the cover to the rear and lift the cover from the system.

Installing the service access cover on a 7063-CR2 system

To install the service access cover on a rack-mounted IBM Power Systems HMC (7063-CR2) system, complete the steps in this procedure.

About this task

⚠ **Attention:** For safety, airflow purposes and thermal performance, the service access cover must be installed and fully seated before you power on the system.

Procedure

1. Ensure that you have the electrostatic discharge (ESD) wrist strap on and that the ESD clip is plugged into a ground jack or connected to an unpainted metal surface. If not, do so now.

2. Place the cover on the system. Align the pins inside the cover with the slots on the top of the chassis as shown in the following figure.

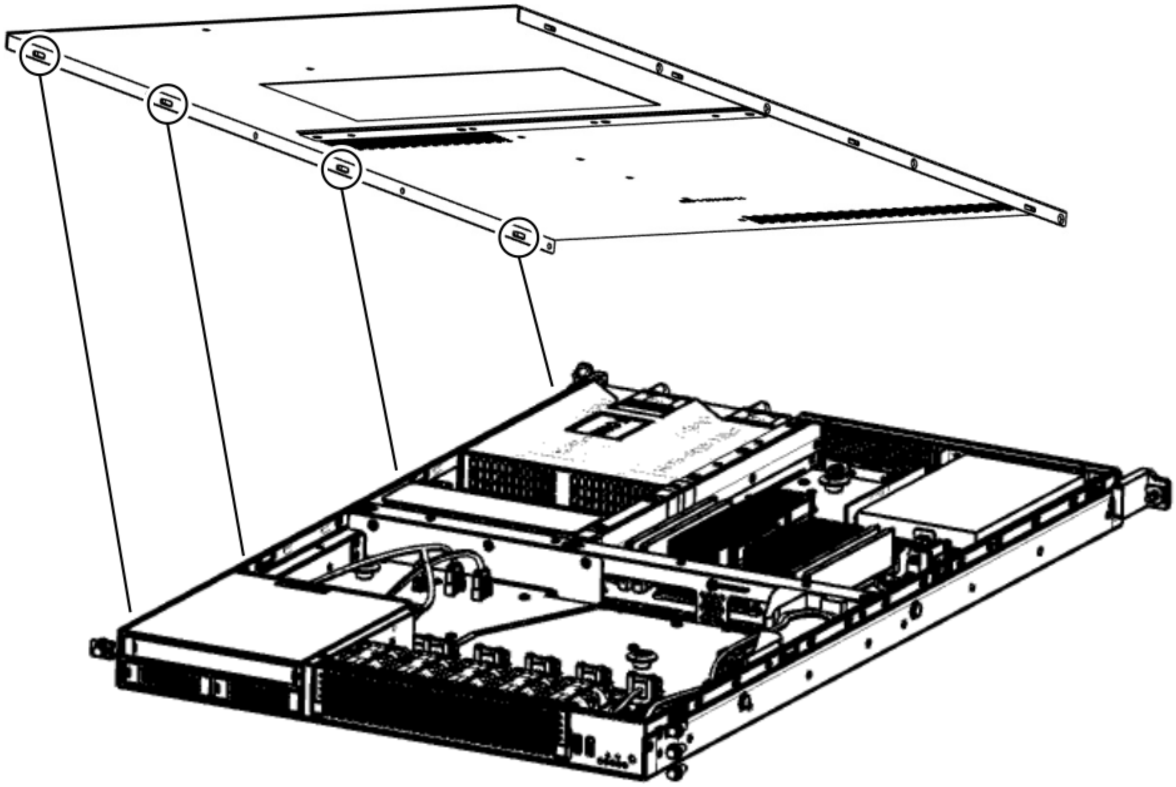

Figure 132. Replacing and securing the cover

3. Slide the cover forwards until it latches into place.
4. Replace the 19 screws to secure the cover as shown in the following figure.

   The system has 4 screws on each side, and 11 screws on the top surface. Use a #2 Phillips screwdriver.

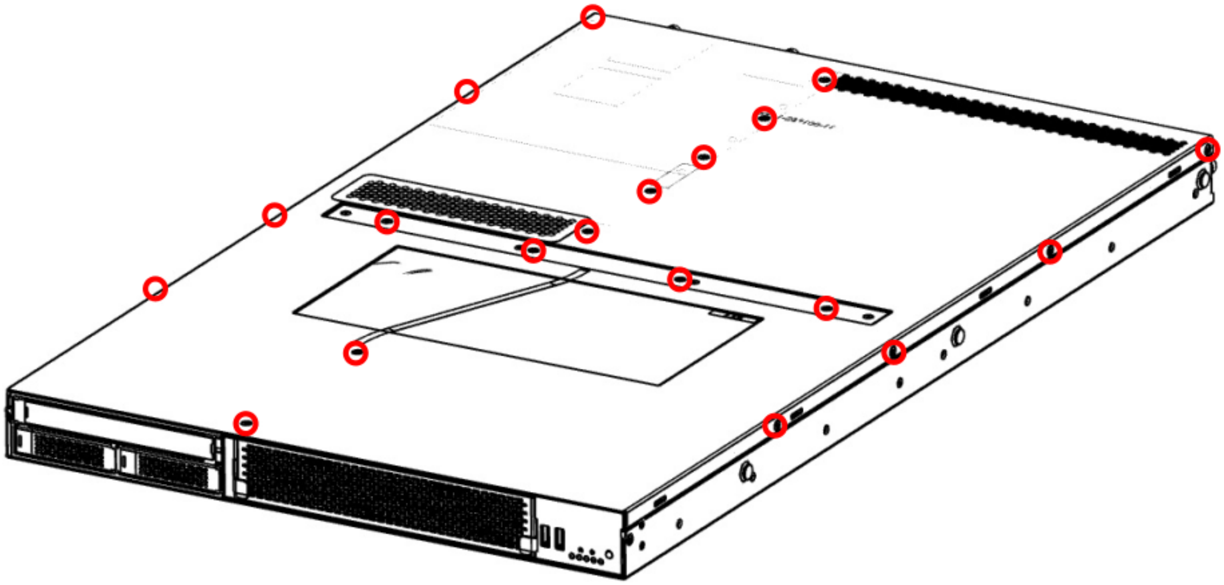

Figure 133. Replacing the cover screws

5. Replace the rails on both sides of the system.
   a) Attach the rail to the system by placing the rail over the three support pins on the side of the system. Ensure that each of the three support pins goes through the rail.

b) Slide the rail forwards onto the support pins.

c) Replace the two screws that secure the rail to the front of the system.

Notices

This information was developed for products and services offered in the US.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with ICT Accessibility 508 Standards and 255 Guidelines (https://www.access-board.gov/ict/) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power servers.

The IBM Power servers online product documentation in IBM Documentation is enabled for accessibility. For more information about IBM's commitment to accessibility, see the IBM accessibility website at IBM Accessibility (https://www.ibm.com/able/).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information.

Electronic emission notices

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Class A Notices

The following Class A statements apply to the IBM servers that contain the Power10 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

The following Class A statements apply to the servers.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany

Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

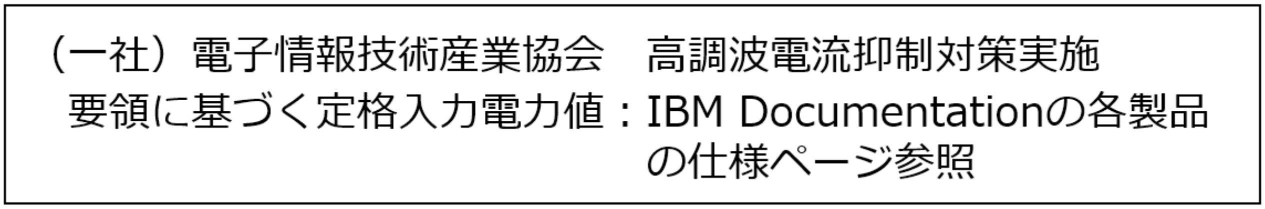

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：IBM Documentationの各製品
の仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

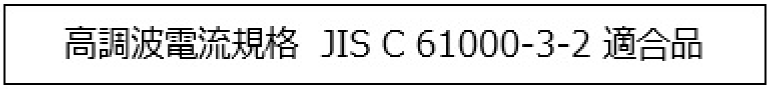

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

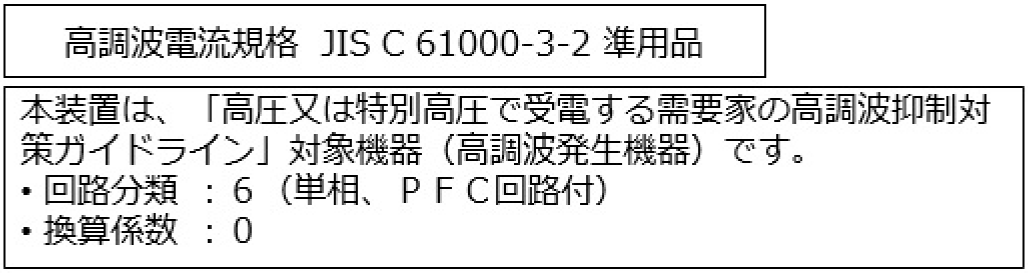

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対
策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

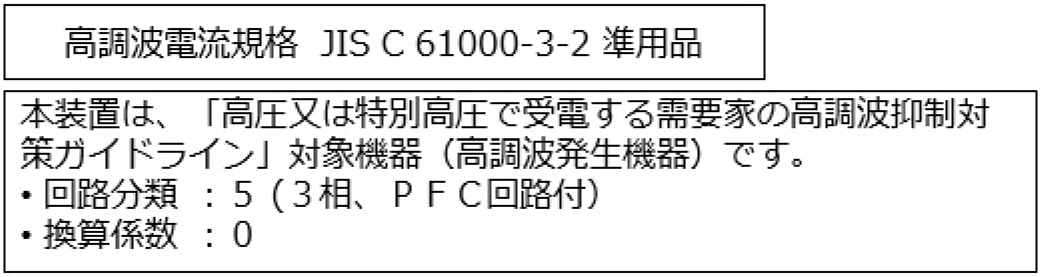

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対
策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

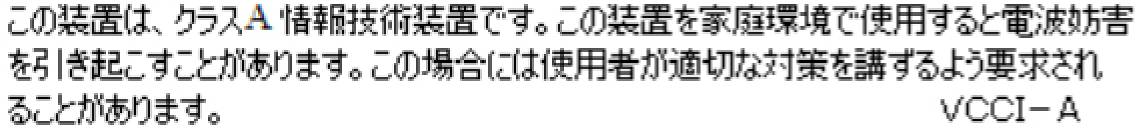

この装置は、クラスA 情報技術装置です。この装置を家庭環境で使用すると電波妨害
を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求され
ることがあります。　VCCI－A

Korea Notice

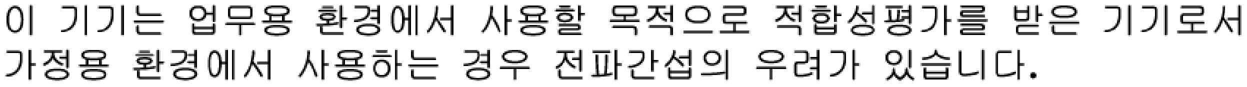

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서
가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

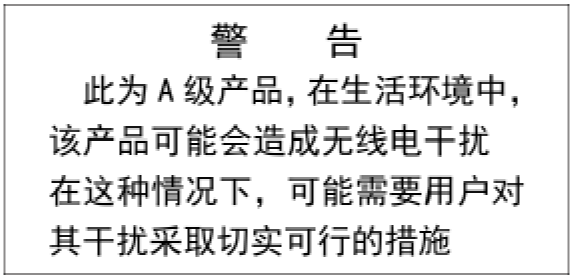

警　告

此为 A 级产品，在生活环境中，
该产品可能会造成无线电干扰
在这种情况下，可能需要用户对
其干扰采取切实可行的措施

Russia Notice

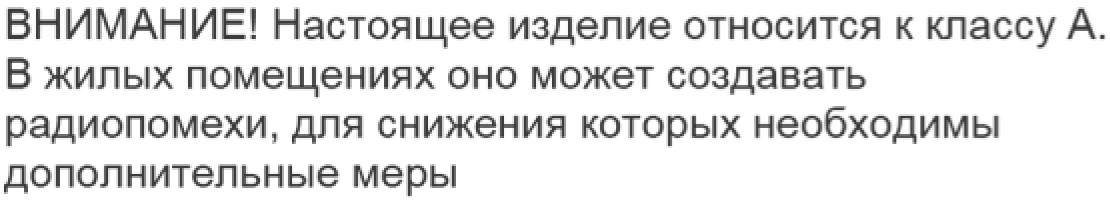

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

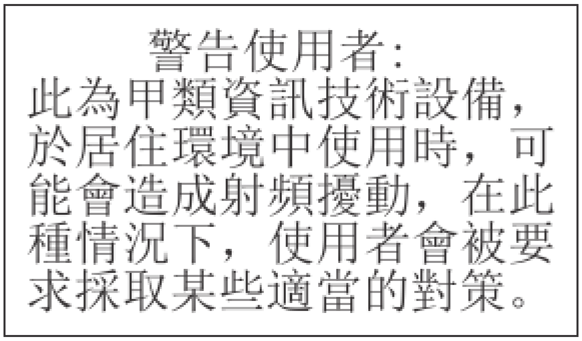

警告使用者：
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

IBM Taiwan Contact Information:

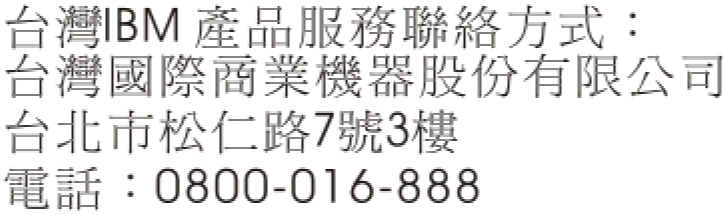

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

United Kingdom Notice

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Canada Notice

CAN ICES-3 (B)/NMB-3(B)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

German Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

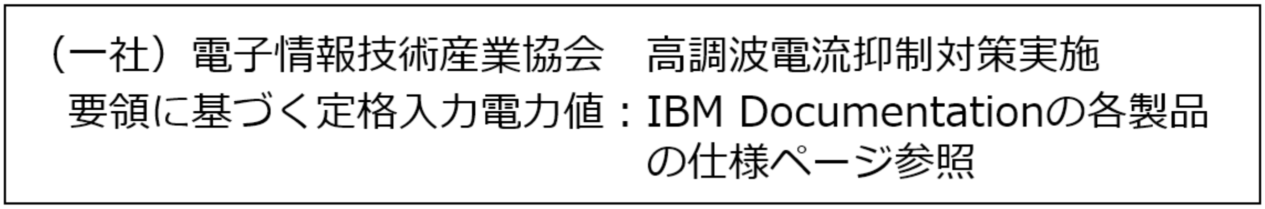

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：IBM Documentationの各製品
の仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

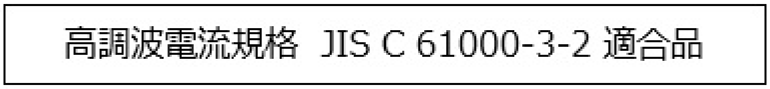

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

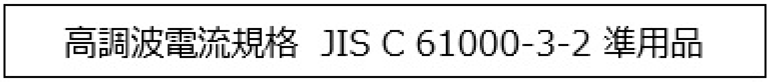

高調波電流規格 JIS C 61000-3-2 準用品

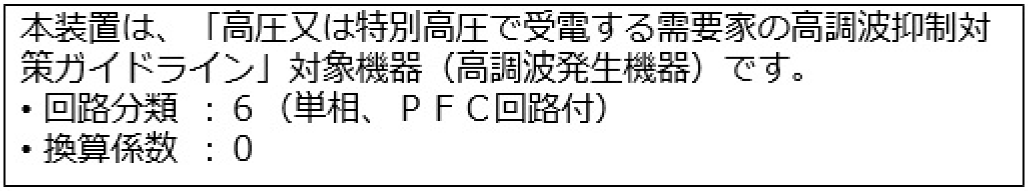

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：６（単相、ＰＦＣ回路付）
- 換算係数 ：０

This statement applies to products greater than 20 A per phase, three-phase.

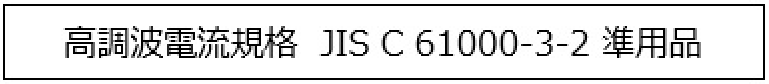

高調波電流規格 JIS C 61000-3-2 準用品

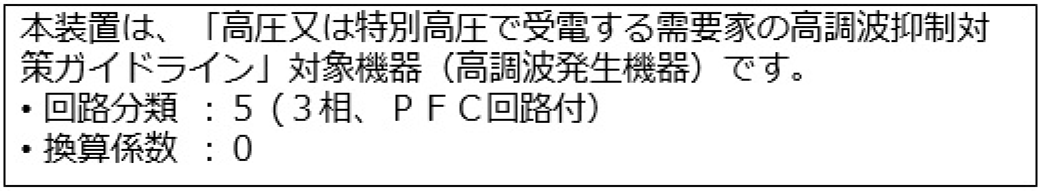

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類 ：５（３相、ＰＦＣ回路付）
- 換算係数 ：０

Japan Voluntary Control Council for Interference (VCCI) Notice

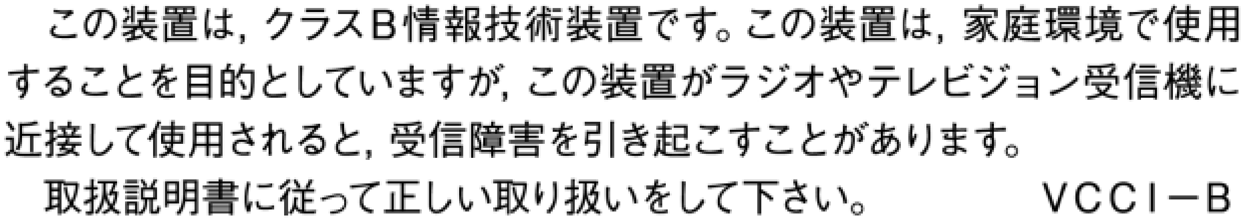

この装置は，クラスB情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。　VCCI－B

Taiwan Notice

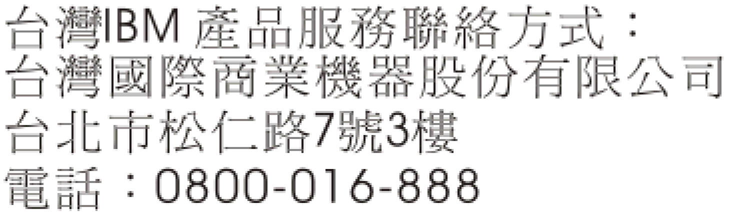

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions:

(1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Responsible Party:

International Business Machines Corporation
New Orchard Road
Armonk, New York 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.