

## IBM i Virtualization and Open Storage (read-me first)

**Mike Schambureck**  
([schambur@us.ibm.com](mailto:schambur@us.ibm.com))

*IBM Lab Services – IBM Power Systems, Rochester, MN*

August 2014



This “read-me first” paper provides detailed instructions on using IBM i 6.1/7.1 virtualization and connecting open storage to IBM i. It provides information on the prerequisites, supported hardware and software, planning considerations, install and post-install tasks such as backups. The paper also contains links to many additional information sources.

1	IBM i virtualization solutions .....	6
1.1	IBM i logical partition (LPAR) hosting another IBM i partition.....	6
1.2	IBM i using open storage as a client of the Virtual I/O Server (VIOS).....	6
1.3	IBM i on a Power blade .....	6
2	IBM i hosting IBM i .....	7
2.1	IBM i hosting IBM i concepts .....	7
2.1.1	VSCSI and Ethernet adapters.....	7
2.1.2	Storage virtualization .....	8
2.1.2.1	Disk virtualization.....	8
2.1.2.2	Tape virtualization .....	9
2.1.2.3	Optical virtualization .....	10
2.1.3	Network virtualization .....	10
2.2	Supported configurations .....	10
2.2.1	Hardware.....	10
2.2.2	Software and firmware.....	11
2.3	Planning for implementing IBM i Client partitions.....	11
2.3.1	Storage planning .....	11
2.3.2	Performance .....	12
2.3.3	Dual/Redundant hosting.....	13
2.3.4	IBM i Client partition considerations and limitations.....	13
2.4	Implementing IBM i client partitions with an IBM i host .....	13
2.4.1	Creating storage spaces for the client lpar .....	13
2.4.2	Creating network server descriptions for the client lpar .....	14
2.4.3	Linking storage spaces to the NWSD .....	14
2.5	Post-install tasks and considerations.....	14
2.5.1	Configure IBM i networking .....	14
2.5.2	How to perform IBM i operator panel functions .....	15
2.5.3	Displaying the IBM i partition System Reference Code (SRC) history .....	15
2.5.4	Configuring Electronic Customer Support (ECS) over LAN .....	15
2.5.5	Copying storage spaces.....	15
2.5.6	Backups.....	16
2.5.7	Operational considerations.....	16
3	IBM i using open storage configurations.....	18
3.1	Storage Area Networks 101 .....	18
3.2	IBM i open storage methods.....	18
3.3	Support for IBM Storage Systems with IBM i .....	19
3.4	Software and firmware: Direct attached .....	20
3.5	Software and firmware: VIOS hosted.....	20
4	IBM i using direct attach storage .....	21
4.1	IBM i SAN configuration steps .....	21
4.1.1	IBM i host kits .....	21
4.1.2	Determine the FC adapter WWPNs.....	21
4.1.3	Configure SAN switch zoning .....	21
4.1.4	IBM i host configuration.....	21
4.1.5	Configure SAN storage for IBM i .....	21
4.1.6	Map the SAN LUNs to the IBM i host .....	22

4.2	Add drives to the IBM i ASP .....	22
5	IBM i hosted by VIOS .....	23
5.1	Virtual SCSI and Storage virtualization concepts .....	23
5.2	Virtual Fibre Channel adapter concepts.....	24
5.2.1	NPIV Supported hardware and software .....	25
5.2.2	Creating Virtual Fibre Channel (VFC) adapters using the HMC .....	25
5.2.3	Mapping VFC adapters to physical ports using the HMC.....	26
5.3	Virtual SCSI (VSCSI) configuration concepts .....	27
5.3.1	VSCSI adapter configuration using IVM .....	28
5.3.2	Creating multiple Virtual SCSI adapters using the VIOS CLI.....	29
5.3.3	Creating multiple Virtual SCSI adapters using the HMC.....	29
5.3.4	Creating multiple Virtual SCSI adapters using the SDMC .....	30
5.3.5	Assigning storage using IVM .....	31
5.3.6	Assigning storage using the HMC .....	31
5.3.7	Assigning storage using SDMC.....	32
5.4	Optical virtualization .....	32
5.4.1	Physical optical devices .....	32
5.4.2	VIOS Media Repository .....	32
5.4.3	Assigning a physical or virtual Optical/DVD/CD drive using the HMC .....	32
5.4.4	Assigning a physical or virtual Optical/DVD/CD drive using the SDMC ...	33
5.4.5	Assigning a physical or virtual Optical/DVD/CD drive using IVM .....	33
5.5	Network virtualization .....	33
5.5.1	Configuring Virtual Ethernet using the HMC .....	33
5.5.2	Configuring Virtual Ethernet using the SDMC .....	34
5.5.3	Configuring Virtual Ethernet using IVM.....	34
5.5.4	Configuring Virtual Ethernet using the VIOS CLI.....	34
5.5.5	Shared Ethernet Adapters (SEAs).....	35
5.5.5.1	Configuring an SEA using IVM .....	35
5.5.5.2	Configuring SEAs using the HMC/SDMC.....	36
5.5.5.3	Configuring SEAs using the VIOS CLI.....	36
5.6	Redundant VIOS Networking.....	36
5.7	Adding new devices to VIOS .....	37
6	Attaching open storage to IBM i through VIOS: Best Practices .....	38
6.1	Storage planning .....	38
6.2	Performance .....	39
6.3	Dual hosting and multi-path I/O (MPIO).....	39
6.3.1	Dual VIOS LPARs with IBM i mirroring.....	39
6.3.2	Path redundancy to a single set of LUNs.....	40
6.3.2.1	Subsystem Device Driver – Path Control Module (SDDPCM).....	40
6.3.2.2	Path redundancy with a single VIOS .....	40
6.3.2.3	Redundant VIOS LPARs with client-side MPIO (VSCSI) .....	41
6.4	Best practices for Fibre Channel storage .....	42
7	Attaching open storage to IBM i through VIOS .....	45
7.1	Open storage configuration.....	45
7.2	VIOS installation and configuration .....	45
7.3	SAN Zoning.....	45

7.4	Mapping storage to IBM i using the HMC .....	46
7.5	Mapping storage to IBM i using the SDMC .....	46
7.6	NPIV Configuration for IBM i LUNs or tape media libraries .....	47
8	IBM i tape options .....	48
8.1	Virtual tape configuration using the HMC .....	48
8.2	Virtual tape configuration using the SDMC .....	49
8.3	Virtual tape configuration using IVM.....	49
9	IBM i installation and configuration .....	50
9.1	Post-IBM i install tasks and considerations .....	50
9.1.1	Configure IBM i networking .....	50
9.1.2	End-to-end LUN device mapping .....	51
9.1.3	How to perform IBM i operator panel functions .....	51
9.1.4	Displaying the IBM i partition System Reference Code history .....	52
9.1.5	Client IBM i LPARs considerations and limitations .....	52
9.1.6	Configuring Electronic Customer Support (ECS) over LAN.....	52
9.1.7	Backups.....	52
10	DS5000 direct attached storage to IBM i.....	53
10.1	Overview.....	53
10.2	Supported hardware and software.....	53
10.3	Best practices, limitations and performance .....	54
10.4	Sizing and configuration .....	55
11	Redundant VIOS Virtual Servers hosting IBM i .....	55
12	Copy Services and IBM i.....	56
12.1	General Support Information .....	56
12.2	DS4000 and DS5000.....	57
12.2.1	FlashCopy and VolumeCopy .....	57
12.2.1.1	Test scenario .....	57
12.2.1.2	FlashCopy and VolumeCopy support statements .....	58
12.2.2	Enhanced Remote Mirroring (ERM) .....	59
12.2.2.1	Test scenario .....	59
12.2.3	ERM support statements .....	60
12.3	SAN Volume Controller (SVC) and Storwize Copy Services and IBM i .....	60
12.3.1	FlashCopy .....	60
12.3.1.1	Test scenario .....	60
12.3.1.2	FlashCopy support statements .....	61
12.3.2	Metro and Global Mirror .....	62
12.3.2.1	Test scenario .....	62
12.3.2.2	Metro Mirror and Global Mirror support statements.....	62
12.4	XIV Copy Services and IBM i.....	63
12.5	DS5000 Direct attach Copy Services.....	63
12.6	NPIV attached storage Copy Services .....	64
13	Appendix:.....	65
13.1	Additional resources .....	65
13.2	IBM i.....	65
13.3	Storage .....	65
13.4	VIOS .....	66

13.5	SVC.....	66
14	Trademarks and disclaimers .....	67

# 1 IBM i virtualization solutions

IBM i 6.1 introduced three significant virtualization capabilities that allow faster deployment of IBM i workloads within a larger heterogeneous IT environment. This section introduces and differentiates these new technologies.

Note: The Oct. 2012 IBM i announcements stated for POWER7+ models: Native IBM i I/O support using IBM i 7.1 is supported, but no native IBM i 6.1 support of I/O. IBM i 6.1 can be a client partition and be provided I/O linkages through either IBM i 7.1 or VIOS.

The Feb. 2013 POWER7+ announcements now provide a way native IBM i 6.1 support of I/O with feature code #EB34 for POWER7+ 710, 720, 730, 740, 770, and 780. There is no native i 6.1 support available at all for POWER7+ 750 and 760.

The May 2013 IBM i announcements now provide native (non-VIOS) attachment of IBM Storwize storage servers. The latest PTFs on IBM i 7.1 TR6 are required for this support.

## 1.1 IBM i logical partition (LPAR) hosting another IBM i partition

An IBM i 6.1/7.1 LPAR can host one or more additional IBM i LPARs, known as virtual client LPARs. Virtual client partitions typically have no physical I/O hardware assigned and instead leverage virtual I/O resources from the host IBM i partition. The types of hardware resources that can be virtualized by the host LPAR are disk, tape, optical and networking. The capability of IBM i to provide virtual I/O resources has been used successfully for several years to integrate AIX®, Linux® and Microsoft® Windows® workloads on the same platform. The same virtualization technology, which is part of the IBM i operating system, can now be used to host IBM i LPARs. IBM i hosting IBM i is the focus of the first half of this paper.

## 1.2 IBM i using open storage as a client of the Virtual I/O Server (VIOS)

IBM i virtual client partitions can also be hosted by VIOS. VIOS is virtualization software that runs in a separate partition with the purpose to provide virtual storage, optical, tape and networking resources to one or more client partitions. The most immediate benefit that VIOS brings to an IBM i client partition is the ability to expand its storage portfolio to use 512-byte/sector open storage. Open storage volumes (or *logical units*, LUNs) are physically attached to VIOS through a FC or a Serial-attached SCSI (SAS) connection and then made available to IBM i. While IBM i does not directly attach to the storage area network (SAN) in this case, as soon as open storage LUNs become available through VIOS, they are managed the same way as integrated disks or LUNs from a directly attached storage system. IBM i using open storage through VIOS is the focus of the second half of this read-me first guide.

## 1.3 IBM i on a Power blade

The third major virtualization enhancement with IBM i 6.1 is the ability to run an IBM i LPAR and its applications on a Power blade server. Running IBM i on a Power blade is beyond the scope of this paper. Refer to the *IBM i on a Power Blade Read-me First* for a complete technical overview and implementation instructions: <http://www.ibm.com/systems/power/hardware/blades/ibmi.html>.

## 2 IBM i hosting IBM i

### 2.1 IBM i hosting IBM i concepts

The capability of an IBM i partition to host another IBM i partition involves hardware and virtualization components. The hardware components are the storage, optical and network adapters and devices physically assigned to the host IBM i LPAR. The virtualization components are the system firmware and IBM i operating system objects necessary to virtualize the physical I/O resources to client partitions. Figure 1 shows the possible solution components:

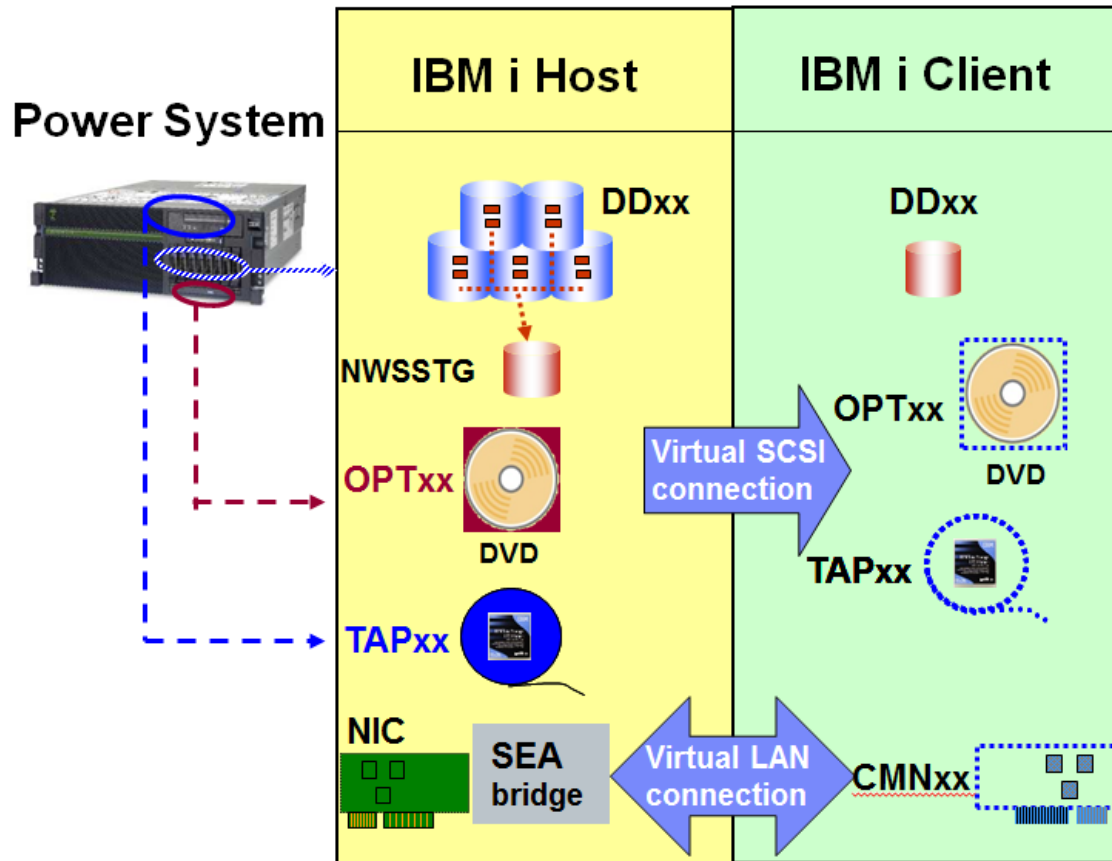


Figure 1: IBM i hosting IBM i components. Note: the IVE/Ethernet adapter represents a bridge that is configured from the virtual Ethernet LAN to a physical Ethernet port. It can also represent an IVE adapter with ports virtualized to the IBM i client partition.

#### 2.1.1 VSCSI and Ethernet adapters

IBM i hosting IBM i uses an existing function of the system firmware, or IBM Power Hypervisor™: the capability to create virtual SCSI (VSCSI) and Ethernet adapters in a partition. Virtual adapters are created for each LPAR in the Hardware Management Console (HMC) or virtual server in the Systems Director Management Console (SDMC). VSCSI adapters are used for storage and optical virtualization; virtual Ethernet adapters are used for network virtualization.

Note that using virtual I/O resources from a host partition does not preclude an IBM i client partition from owning physical hardware. A mix of virtual and physical hardware in the same

partition is supported for IBM i in this environment, by assigning both the types of adapters to the partition in the HMC or SDMC.

## 2.1.2 Storage virtualization

### 2.1.2.1 Disk virtualization

To virtualize integrated disk (SCSI, SAS or SSD) or LUNs from a SAN system to an IBM i client partition or virtual server, both HMC/SDMC and IBM i objects must be created. In the HMC/SDMC, the minimum required configuration is:

- One VSCSI server adapter in the host partition
- One VSCSI client adapter in the client partition

Note: for the remainder of this paper, the terms LPAR or partition (when using an HMC) will also mean virtual server (when using SDMC).

This VSCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one VSCSI pair can exist for the same client partition in this environment.

There is no additional disk storage configuration required in the client partition. In the host partition, the minimum required setup consists of the following requirements:

- One network server description (NWS) object
- One network server storage space (NWSSTG) object

At least one NWS object must be created in the hosting partition for each client partition. The NWS object is associated with a VSCSI server adapter in the hosting partition's profile, which in turn is connected to a VSCSI client adapter in the HMC/SDMC. The NWS also contains links to one or more NWSSTG objects. See section "Creating network server descriptions for the client lpar" for command details.

The NWSSTG objects correspond to the virtual disks provided to the client IBM i partition. They are created from available physical storage in the host partition. In the client, they are recognized and managed as standard DDxx disk devices (with a different type and model). Figure 2 shows several storage spaces for a client partition in an IBM i 6.1/7.1 host partition:

```

Work with Network Server Storage Spaces
System: ZA6BP3
Type options, press Enter.
 1=Create  2=Change  3=Copy  4=Delete  5=Display  6=Print  10=Add link
11=Remove link

  Opt  Name      Server  Seq  Link  Access  Stg
  ---  ---      ---    ---  ---   ---     ---
  ---  CP10D01   CP10A   2    *DYN *UPDATE
  ---  CP10D02   CP10A   3    *DYN *UPDATE
  ---  CP10D03   CP10A   4    *DYN *UPDATE
  ---  CP10D04   CP10A   5    *DYN *UPDATE
  
```

Figure 2: example storage spaces

Figure 3 shows several storage spaces as seen in an IBM i 6.1/7.1 client partition:



```

Work with Disk Status
Y213B2P2
03/26/08 16:43:00
Elapsed time: 00:00:00

```

Unit	Type	Size (M)	% Used	I/O Rqs	Request Size (K)	Read Rqs	Write Rqs	Read (K)	Write (K)	% Busy
1	6B22	34359	14.6	.0	.0	.0	.0	.0	.0	0
2	6B22	34359	5.6	.0	.0	.0	.0	.0	.0	0
3	6B22	34359	5.6	.0	.0	.0	.0	.0	.0	0
4	6B22	34359	5.6	.0	.0	.0	.0	.0	.0	0
5	6B22	34359	5.6	.0	.0	.0	.0	.0	.0	0

Figure 3: Example of what the client partition sees as storage.

Storage spaces for an IBM i client partition do not have to match physical disk sizes; they can be created from 160 MB to 1 TB in size, as long as there is available storage in the host. The 160 MB minimum size is a requirement from the storage management Licensed Internal Code (LIC) on the client partition. NOTE: To anticipate i7.2 load source size enforcements, you should use a storage space size of at least 80GB.

An **IBM i client** partition supports up to 16 NWSSTGs (i7.1 TR8 increased this limit to 32) linked to a single NWSD, and therefore, to a single VSCSI connection. If more storage spaces are required, you need to create another set of VSCSI adapters on the HMC and another NWSD associated with the same client LPAR on the IBM i. Up to 32 outstanding I/O operations from the client to each storage space are supported for IBM i clients. Storage spaces can be created in any existing auxiliary storage pool (ASP) on the host, including Independent ASPs. Through the use of NWSSTGs, any physical storage supported in the IBM i host partition can be used to provide virtual disk storage to a client partition.

For performance reasons, consider creating multiple storage spaces associated with multiple NWSDs. The rule of thumb is six to eight storage spaces for each client partition at a minimum. This implies that you are also creating multiple sets of VSCSI adapter pairs between the hosting partition and the client partition. Associate each hosting partition's server VSCSI adapter with a separate NWSD by referencing the VSCSI adapter's resource name in the NWSD. Then link storage spaces to the NWSDs. This will supply multiple disk arms for the client partition to use. See sections "Creating storage spaces for the client lpar" and "Linking storage spaces to the NWSD" for command details>

### 2.1.2.2 Tape virtualization

Specific IOP-less attached physical tape devices supported in the host IBM i LPAR can be virtualized to an IBM i client LPAR. See the links below for specific details. An existing VSCSI connection can be used, or a new connection can be created explicitly for tape I/O traffic. By default, if a VSCSI connection exists between the host and the client, all TAPxx drives supported for virtualization by the host will be available to the client, where they will also be recognized as TAPxx devices. The NWSD parameter **Restricted device resources** can be used to specify which tape devices in the host a client partition **cannot** access.

A virtualized tape drive in the host partition can be used for a D-mode initial program load (IPL) and for the install of the client partition or its applications. The client partition will be able to write to the physical media in the drive. Although Virtual tape devices backed by image catalogs (TAPVRTxx) cannot be virtualized to an IBM i client. These restrictions are documented in the IBM i InfoCenter at:

<http://publib.boulder.ibm.com/infocenter/iseres/v7r1m0/index.jsp?topic=%2Fzram4%2Fzram4clie%2Fvirtdevices.htm>

For the current list of supported tape libraries and associated operating system levels please refer to the DeveloperWorks web page for IBM Removable Media on IBM i:

### 2.1.2.3 Optical virtualization

Any optical drive supported in the host IBM i LPAR can be virtualized to an IBM i client LPAR. An existing VSCSI connection can be used, or a new connection can be created explicitly for optical I/O traffic. By default, if a VSCSI connection exists between host and client, all physical and virtual OPTxx optical drives in the host will be available to the client, where they will also be recognized as OPTxx devices. The NWSD parameter **Restricted device resources** can be used to specify which optical devices in the host a client partition **cannot** access.

A virtualized optical drive in the host partition can be used for a D-mode Initial Program Load (IPL) and installation of the client partition, as well as for installing Program Temporary Fixes (PTFs) or data. If the optical drive is writeable, the client partition will be able to write to the physical media in the drive.

### 2.1.3 Network virtualization

Using a virtual LAN (VLAN) for partition-to-partition communication within a system and routing Ethernet traffic between a virtual LAN and an external physical LAN are existing IBM i capabilities. In order to route a client partition's virtual Ethernet traffic through a host's physical network adapter, a virtual Ethernet adapter must be created in the HMC in both partitions or the SDMC in both virtual servers. To be on the same VLAN, the two virtual Ethernet adapters must have the same Port Virtual LAN ID (PVID). This type of adapter is recognized by IBM i as a communications port (CMNxx) with a different type (268C). In the host partition that is at IBM i 7.1 TR3 level or newer, the virtual Ethernet adapter is then associated with the physical network adapter via a **BRIDGE** parameter configuration on the line descriptions associated with both the physical and virtual Ethernet adapters. This allows the client partition to send network packets via the VLAN out the physical adapter port to the outside LAN. The physical adapter can be any network adapter supported by IBM i 6.1/i7.1. An Integrated Virtual Ethernet (IVE) port, also known as Host Ethernet Adapter (HEA) ports can NOT be bridged to an external LAN. These ports are managed by the Power server's hypervisor and are allocated directly to the client partition without the involvement of VIOS or virtual Ethernet adapters.

## 2.2 Supported configurations

### 2.2.1 Hardware

One of the most significant benefits of this solution is the broad hardware support. Any storage and network adapters and devices supported by the host IBM i partition on an IBM POWER6® or IBM POWER7® processor-based server can be used to provide virtual I/O resources to the client IBM i partition. Virtualization of tape devices from an IBM i host to an IBM i client is supported as of June 2011. Table 1 lists the hardware supported.

Hardware type	Supported for IBM i hosting IBM i	Notes
IBM Power servers	Yes	<b>Includes</b> IBM POWER6 AND POWER7 servers  <b>Does not include</b> IBM POWER6/7 processor-based blades.
IBM Power 575/775 servers	No	

IBM POWER5-based systems or earlier	No	
Storage adapters (FC, SAS, SCSI)	Yes	Must be supported by IBM i 6.1 or higher <b>and</b> supported on POWER6/7-based IBM Power server
Storage devices and subsystems	Yes	Must be supported by IBM i 6.1 or higher <b>and</b> supported on POWER6/7-based IBM Power server
Network adapters	Yes	Must be supported by IBM i 6.1 or higher <b>and</b> supported on POWER6/7-based IBM Power server
Optical devices	Yes	Must be supported by IBM i 6.1 or higher <b>and</b> supported on POWER6/7-based IBM Power server
Tape devices	Yes*	Physical tape devices are supported when the hosting IBM i is at i7.1 TR2 or newer.

Table 1: Hardware supported

To determine the storage, network and optical devices supported on each IBM Power server model, refer to the sales manual for each model at: <http://www.ibm.com/common/ssi/index.wss>. To determine the storage, network and optical devices supported only by IBM i 6.1, refer to the upgrade planning website at: <https://www-304.ibm.com/systems/support/i/planning/upgrade/futurehdwr.html>.

## 2.2.2 Software and firmware

Software or firmware type	Supported for IBM i hosting IBM i	Notes
IBM i 6.1 or higher	Yes	<b>Required</b> on both host and client IBM i partition
IBM i 5.4 or earlier	No	Not supported on host or client partition
IBM Power server system firmware 320_040_031 or higher	Yes	This is the minimum system firmware level required
HMC firmware V7 R3.2.0 or higher	Yes	This is the minimum HMC firmware level required
SDMC firmware V6 R7.3.0 or higher	Yes	This is the minimum HMC firmware level required

Table 2: Software and firmware requirements

## 2.3 Planning for implementing IBM i Client partitions

### 2.3.1 Storage planning

Because the storage for virtual disks in the IBM i client LPAR is provided by NWSSTG objects in the host LPAR, the main prerequisite to installing a new client LPAR is having sufficient capacity in the host to create those objects. Note that the host partition is not capable of detecting the percent of the virtual storage that is used by the client. For example, if a 500-GB storage space is created, it occupies that amount of physical storage in the host IBM i LPAR, even if the disk capacity is only 50% utilized in the client LPAR. The number of NWSSTG objects is important for the client IBM i partition to allow for concurrent disk IO.

It is recommended to closely match the total size of the storage spaces for each client partition to its initial disk requirements. As the storage needs of the client partition grow, additional storage spaces can be dynamically created and linked to it on the host partition. On the client, the new virtual disk will automatically be recognized as a non-configured drive and can be added to any existing ASP. The only restriction to consider in this case is the maximum number of storage spaces allowed per VSCSI connection for an **IBM i client** partition, which is 16 (i7.1 TR8 increased this limit to 32). If more than 16/32 NWSSTGs are needed for a client LPAR, additional VSCSI connections can be created dynamically in the HMC or the SDMC. Repeat the steps used for defining the first VSCSI adapter pair and associating the server adapter on the hosting partition with another network server description (NWSN). On the NWSN, set the partition name or number to the same client partition you used for the first NWSN. Always set the power control parameter of any additional NWSNs to \*NO.

### 2.3.2 Performance

As described in the “IBM i logical partition (LPAR) hosting another IBM i partition” section, disk I/O operations in an IBM i virtual client partition result in I/O requests to the physical disk adapters and drives that are assigned to the host partition. Therefore, *the best way to ensure good disk performance in the client LPAR is to create a well-performing disk configuration in the host LPAR*. As the host partition is a standard IBM i partition, all the recommendations in the *Performance Capabilities Reference* manual at:

<http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>) will apply to it. Use the suggestions provided in the manual for maximizing IBM i disk performance for the type of physical storage used in the host, whether it is integrated disk or SAN.

Note that if only the system ASP exists on the host partition, NWSSTG objects are created on the same physical disk units as all other objects. If the host partition is running production applications in addition to providing virtual storage to client partitions, there will be disk I/O contention as both client partitions and IBM i workloads in the host send I/O requests to those disk units. To minimize disk I/O contention, create storage space objects in a separate ASP on the host (Independent ASPs are supported). Performance on the client(s) would then depend on the disk adapter and disk configuration used for that ASP. If the host partition is providing virtual storage to more than one client partition, consider using separate ASPs for the storage space objects for each client. This recommendation should be weighed against the concern of ending up with too few physical disk arms in each ASP to provide good performance.

Disk contention from IBM i workloads in the host LPAR and virtual client LPARs can be eliminated if a separate IBM i LPAR is used just for hosting client LPARs. An additional benefit of this configuration is the fact that an application or OS problem stemming from a different workload on the host cannot negatively affect client partitions. These benefits should be weighed against:

- The license cost (per core) associated with a separate IBM i partition
- The maintenance time required for another partition, such as applying Program Temporary Fixes (PTFs)
- The ability to create well-performing physical disk configurations in both partitions that meet the requirements of their workloads

If the host partition runs a heavy-I/O workload and the client partitions also have high disk response requirements, it is strongly recommended to consider using a separate hosting partition, unless separate ASPs on the host are used for storage space objects. If the host partition’s workload ranges from light to moderate with respect to disk requirements and the client partitions are used mostly for development, test or quality assurance (QA), it is acceptable to use one IBM i partition for both tasks.

The hosting IBM i partition will experience an increase in the paging of its memory pages as a result of its hosting role. Monitor and adjust the memory pools using the tips in section “General

Performance Information, Tips, and Techniques” of the *Performance Capabilities Reference* manual at: (<http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>)

### **2.3.3 Dual/Redundant hosting**

An IBM i client partition has a dependency on its host: if the host partition fails, IBM i on the client will lose contact with its disk units. The virtual Ethernet connection would also become unavailable if the host partition is brought down to restricted state or shut down for scheduled maintenance or to apply PTFs. To remove this dependency, two host partitions can be used to simultaneously provide virtual storage to one or more client partitions.

The configuration for two hosts for the same client partition uses the same concepts as that for a single host described in the “IBM i logical partition (LPAR) hosting another IBM i partition” section. In addition, a second VSCSI client adapter exists in the client LPAR, connected to a VSCSI server adapter in the second host LPAR. The IBM i configuration of the second host mimics that of the first host, with the same number of NWSD and NWSSTG objects, and NWSSTG objects of the same size. As a result, the client partition recognizes a second set of virtual disks of the same number and size. To achieve redundancy, adapter-level mirroring is used between the two sets of storage spaces from the two hosts. Thus, if a host partition fails or is taken down for maintenance, mirroring will be suspended, but the client partition will continue to operate. When the inactive host is either recovered or restarted, mirroring can be resumed.

### **2.3.4 IBM i Client partition considerations and limitations**

Refer to the topic *Considerations and limitations for i5/OS client partitions on systems managed by the Integrated Virtualization Manager (IVM)* in the Information Center at: <http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahc/rzahcbladei5limits.htm>. While in this case the IBM i client partition is not being managed by IVM, it does virtual I/O resources and the limitations outlined in the topic above apply to it.

## **2.4 Implementing IBM i client partitions with an IBM i host**

Installing IBM i in a client LPAR with an IBM i host consists of two main phases:

- Creating the VSCSI configuration in the HMC or SDMC
- Creating the NWSSTG and NWSD objects in the IBM i host partition, and activating the new client partition

The implementation steps for creating VSCSI adapter pairs are described in detail in the topic *Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC* in the *Power Systems Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>. Note that sufficient available capacity is required in the IBM i host partition to create the storage space objects. When following the detailed implementation instructions, keep in mind the performance recommendations explained in the “Performance section of this paper.

### **2.4.1 Creating storage spaces for the client lpar**

You can use the CRTNWSSTG command to create individual storage spaces in the hosting partition. They default to using storage from the system ASP (#1) but can also use user or independent ASPs (ie #2-33).

These commands will take a while to run, so it is recommended that you submit them to a batch job running in the background and possibly off-shift. Here is an example command:

- SBMJOB CMD(CRTNWSSTG NWSSTG(TESTLPAR01) NWSSIZE(102400) FORMAT(\*OPEN) OFFSET(\*ALIGNLGLDSK) TEXT('TESTLPAR - storage space 01')) JOB(TESTLPAR01)
  - This command will be submitted as a batch job called TESTLPAR01.
  - It will create a 100 GB storage space called TESTLPAR01 with the proper parameters for a client LPAR.
  - Repeat the command, incrementing the storage space number, for each of the required storage spaces.

## 2.4.2 Creating network server descriptions for the client lpar

Creating the VSCSI server adapters in the hosting IBM i partition will generate **CTLXX** communications resource names in the hosting server. Use the **WRKHDWRSC \*CMN** command to see these resources. Take an option 7 for details and look at the adapter *slot number* on the hardware resource. This will match the *adapter ID* on the VSCSI adapter from the HMC. Using the CTLxx resource name in the creation of the network server description (NWSD) will associate the VSCSI adapter to this server description. The *partition ID* or the *partition name* are also specified in the NWSD to further tie the pieces together. Here is an example of the command:

- CRTNWSW NWSD(TESTLPAR) RSRNAME(CTL01) TYPE(\*GUEST \*OPSYS) PARTITION(TESTLPAR) IPL SRC(\*PANEL) PWRCTL(\*NO) TEXT('NWSD presents storage for the client testlpar')

## 2.4.3 Linking storage spaces to the NWSD

Finally, you need to map or link the storage spaces to the NWSD so that the storage will be virtualized to the client partition. Here is an example of the command:

- ADDNWSSTGL NWSSTG(TESTLPAR01) NWSD(TESTLPAR)
  - Repeat the command for up to 16 storage spaces (32 if using i7.1 TR8 or newer).

## 2.5 Post-install tasks and considerations

### 2.5.1 Configure IBM i networking

Once the IBM i client partition is installed and running, the first system management step is to configure networking. There are three types of network adapters that can be assigned to an IBM i client partition:

- A standard physical network adapter in a Peripheral Component Interconnect (PCI) slot
- A logical port on a HEA
- A virtual Ethernet adapter

Note that both physical and virtual I/O resources can be assigned to an IBM i virtual client partition. If a physical network adapter was not assigned to the IBM i client partition when it was first created, refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Power Systems Logical Partitioning Guide* at: (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>) to assign an available adapter.

An IBM i client partition can also use the HEA capability of POWER6 processor-based servers. To assign a logical port (logical HEA or LHEA) on an HEA to an IBM i client partition, refer to the topic **Creating a Logical Host Ethernet Adapter for a running logical partition using the HMC** in the *Power Systems Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.

Lastly, a virtual Ethernet adapter can also provide network connectivity to an IBM i client partition. To create one, refer to the topic **Configuring a virtual Ethernet adapter using the HMC** in the *Power Systems Logical Partitioning Guide* at:  
<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.

In all three cases, the assigned network adapter will be recognized as a communications port (CMNxx) in IBM i. The type of communications port will depend on the network adapter: a Gigabit Ethernet adapter port, a logical HEA or a virtual Ethernet adapter. In the case of a standard PCI network adapter or an logical HEA, networking can be configured following the process described in the IBM i networking topic in the Information Center at:  
<http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm>.

If the IBM i client partition is to use a virtual Ethernet adapter to communicate with an external network, additional configuration on the IBM i host is required. The virtual Ethernet adapter allows the client partition to communicate only with other partitions whose virtual Ethernet adapters have the same Port Virtual LAN ID (PVID); in other words, partitions on the same virtual LAN within the system. A routing configuration can be created in the IBM i host partition to allow forwarding of network packets from the outside LAN to the client partition on the virtual LAN. With i 7.1 TR3 PTFs installed on the hosting IBM i partition, a bridge can be configured between the virtual and physical Ethernet adapters. See the steps to do this in section **Ethernet Layer-2 Bridging** of this guide: <http://www.redbooks.ibm.com/abstracts/redp4806.html?Open>

## 2.5.2 How to perform IBM i operator panel functions

Operator panel functions in an IBM i client partition are performed in the HMC:

- Sign onto the HMC with a profile that has sufficient authority to manage the IBM i client partition
- Select the partition
- Use the open-in-context arrow to select **Serviceability → Control Panel Functions**, then the desired function.

## 2.5.3 Displaying the IBM i partition System Reference Code (SRC) history

- Sign onto the HMC with a profile that has sufficient authority to manage the IBM i client partition.
- Select the partition.
- Use the open-in-context arrow to select **Serviceability → Reference Code History**.
- To display words 2 through 9 of a reference code, select the option specific to that code.

## 2.5.4 Configuring Electronic Customer Support (ECS) over LAN

A supported WAN adapter can be assigned to the IBM i client partition for ECS. Alternatively, ECS over LAN can be configured. Refer to the topic **Setting up a connection to IBM** in the Information Center at:  
[http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzaji/rzaji\\_setup.htm](http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm).

## 2.5.5 Copying storage spaces

Because an IBM i client partition is installed into one or more storage space objects in the IBM i host partition, new client partitions can be deployed rapidly by copying the storage space(s). Note that each IBM i partition, client or host, must have valid OS and Licensed Program Product (LPP) licenses for the number of processors it uses.

To copy one or more storage spaces that contain an installed IBM i client partition, first shut down the partition during an available period of downtime. Next, log into the host IBM i partition with a security officer-level profile and perform the following steps:

- Enter **WRKNWSSTG**.
- Enter **3** next to the storage space you are going to copy and press **Enter**.
- Enter a name that can be a maximum of 10 characters for the new storage space.
- The size of the original storage space will be entered automatically. The new storage space can be as large or larger (a maximum of 1 TB), but not smaller.
- Enter the correct ASP ID. The ASP where the original storage space exists is the default
- Optionally, enter a text description and press **Enter**.

To deploy the new client partition, follow the instructions in the “Implementing IBM i client partitions with an IBM i host” section to create the necessary VSCSI configuration in the HMC and the NWSD object in the host IBM i partition.

## 2.5.6 Backups

As mentioned above, an IBM i client partition with an IBM i host can use a mix of virtual and physical I/O resources.

If the hosting partition has the IBM i7.1 TR2 PTFs applied, specific physical tape devices attached to the hosting partition can be virtualized to client partition(s) with the proper 6.1.1 or 7.1 PTFs applied.

Another approach is to assign a physical tape adapter on the system to it and treat it as a standard IBM i partition. The tape adapter can be any adapter supported by IBM i on IBM Power servers and can be shared with other partitions.

To assign a physical tape adapter to the IBM i client partition, refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>.

When a tape adapter connected to a tape drive or library is available to the client partition, refer to the Backup and Recovery topic in the Information Center to manage backups at: [http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int\\_215989](http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int_215989).

The IBM i host partition can also be used for **system-level** backups of the client partition. Refer to the topic **Saving IBM i server objects in IBM i** in the *Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>.

## 2.5.7 Operational considerations

Remember that the IBM i host partition's state may affect the IBM i client partition. For instance, when the hosting partition is restarted, the client partition loses contact with its disk. This results in a system reference code sequence of A6xx0255 shown for the client partition. The client VSCSI adapter loses its virtual path to the server when the server is rebooted and any I/O that was in flight is returned to the device drivers as "aborted". The client's disk device driver will retry the operation, and receive failure responses while the server is still unavailable. Eventually the client's device driver gives up and returns an unsuccessful response to the microcode's Storage Management code. That microcode puts up the 0255 SRC and retries the aborted operations indefinitely until the disks are reconnected. There will be no data loss/data corruption as long as the client partition remains active until the disk path is reconnected.



As for Ethernet traffic in this same scenario, the TCP subsystem in the hosting partition would come down, so the virtual networking in the client partition would no longer have a job to communicate with. So the client partition would lose communications during this scenario. The management console supplied IBM i client console would need to be used to restart the communications.

If the hosting partition is put into restricted state, the Ethernet communications for the client partition would stop, but the disk activity, say for a running batch job, would continue to run. The Network Server Description (NWSD) in the hosting partition is not automatically varied off and is not part of an ended subsystem, so the storage spaces associated with the NWSD remain available to the client partition.

## 3 IBM i using open storage configurations

### 3.1 Storage Area Networks 101

There are some basic SAN concepts that need to be understood by a person coming from an integrated disks background:

- The disk drives installed in a SAN are not directly seen by IBM i. They are typically grouped into RAID sets of drives called arrays, extent pools or storage pools. RAID 0, 1, 5, 6 and 10 are generally supported. The disks are RAIDed in the SAN and should not\* be protected on IBM i.
  - \* Note, the DS8000 has an option to present unprotected LUNs directly to IBM i, where they would be protected by IBM i.
  - Most SANs support different types of disk drives (ie fibre channel versus SATA) and different disk drive speeds. Typical IBM i workloads work best with faster FC or SAS drives versus SATA drives.
- After the RAID array is created, the SAN interface is used to create logical drives from the space within the array.
  - The user sets the size of the logical drives from a range of a few GBs to almost 2 TBs each. IBM i can handle drives anywhere in this range, except for restrictions on the load source drive (70 GB minimum, 1.9 TB maximum). There is no need to match the size of the supported integrated (internal) disk drives.
  - IBM i performs best with logical drives that are the same size.
  - IBM i has always been architected to perform best with more disk arms. This does not change with SAN disks. You need to create a good number of logical drives, not one large drive. Refer to the Storage planning section for more details.
- The next step is to map or assign the logical drives to a host server. This server can be VIOS or a direct connection to IBM i (depending upon the SAN being used. Refer to the next two sections for details).
  - Hosts are known by the Fibre Channel World Wide Port Name (WWPN) or the SAS MAC address of the port on the adapter that the host server owns.
  - A separate step, called SAN zoning, has to take place to configure the SAN switches to connect the host adapter ports to the SAN's host ports.
  - As the logical drives are mapped to the host they are assigned a logical unit number or LUN. A LUN is a unique number within an array. Many people refer the logical drive as a LUN.
- On the VIOS host the mapped logical drives are seen as hdisks and then assigned to an IBM i partition. It is recommended to assign whole hdisks to IBM i. VIOS supports logical volumes where you can subset an hdisk into multiple volumes to assign to a partition. This is not recommended for IBM i partitions.
- On an IBM i, the VIOS assigned drives or the direct attached mapped logical drives are seen as DDxxx drives. The drives can be shown as "unprotected" but the SAN is providing the protection. The drives can be added to any auxiliary storage pool (ASP) that you choose.

### 3.2 IBM i open storage methods

There are three general methods by which IBM i connects to open storage:

- Directly through the SAN fabric without VIOS.
- As a client of VIOS, with VIOS connecting to a storage subsystem through the SAN fabric. In this case VIOS sees the LUNs and then maps them as virtual SCSI (VSCSI) LUNs to the IBM i client partition.

- As a client of VIOS, with VIOS facilitating the access to the storage using N-Port ID virtualization (NPIV). In this case VIOS does not see the LUNs, they are mapped to a virtual WWPN that the IBM i partition owns. More on that later

The table below applies to both of the connection methods: IBM i direct attachment and IBM i → VIOS → storage subsystem. This paper does not attempt to list the full device support of VIOS, nor of any other clients of VIOS, such as AIX and Linux. For the general VIOS support statements, including other clients, refer to the VIOS datasheet at: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>.

IBM i as a client of VIOS is also supported on POWER6 or higher processor-based blade servers. IBM i cannot attach directly to open storage when running on Power blades. For the full supported configurations statement for the IBM i on Power blade solution, refer to the [BladeCenter Interoperability Guide \(BIG\)](http://www-947.ibm.com/support/entry/portal/docdisplay?Indocid=MIGR-5073016) at: <http://www-947.ibm.com/support/entry/portal/docdisplay?Indocid=MIGR-5073016>

### 3.3 Support for IBM Storage Systems with IBM i

As of April, 2014 this chart shows the supported IBM storage systems with IBM i, both direct attachment and hosted attachment.

Table as of April, 2014		DS3200 DS3400 DS3500 DS3700 DS3950	DS4700 DS4800 DS5020	SVC Storwize V7000 V5000 V3700 V3500	DS5100 DS5300	XIV	DS8100 DS8300 DS8700 DS8800 DS8870
Rack / Tower Systems	IBM i Version Hardware	6.1 / 7.1 POWER6/7 Not DS3200#, Yes DS3500##	6.1 / 7.1 POWER6/7	6.1 / 7.1 POWER6/7	6.1 / 7.1 POWER6/7	6.1 / 7.1 POWER6/7	5.4 / 6.1 / 7.1 POWER5/6/7
	IBM i Attach	VIOS VSCSI	VIOS VSCSI	Direct*** or VIOS VSCSI and NPIV%%	Direct* or VIOS -VSCSI and NPIV%	VIOS VSCSI	Direct or VIOS VSCSI and NPIV**
Power Blades or Compute Nodes	IBM i Version Hardware	6.1 / 7.1 POWER6/7 @, #, ##	6.1 / 7.1 POWER6/7 (BCH)	6.1 / 7.1 POWER6/7 (BCH)	6.1 / 7.1 POWER6/7 (BCH)	6.1 / 7.1 POWER6/7 (BCH)	6.1 / 7.1 POWER6/7 (BCH)
	IBM i Attach	VIOS VSCSI	VIOS VSCSI	VIOS VSCSI NPIV%%	VIOS VSCSI and NPIV%	VIOS VSCSI	VIOS VSCSI and NPIV**

Notes:

- This table does not list more detailed considerations, for example required levels of firmware or PTFs required or configuration performance considerations

- POWER7 servers require IBM i 6.1 or later

# DS3200 only supports SAS connection, not supported on Rack/Tower servers which use only Fibre Channel connections, supported on Blades with SAS

## DS3500 has either SAS or Fibre Channel connection. Rack/Tower only uses Fibre Channel. Blades in BCH support either SAS or Fibre Channel. Blades in BCS only uses SAS.

### Not supported on IBM i 7.1. But see SCORE System RPQ 846-15284 for exception support

\* Supported with Smart Fibre Channel adapters – NOT supported with IOP-based Fibre Channel adapters

\*\* NPIV requires Machine Code Level of 6.1.1 or later and requires NPIV capable HBAs (FC adapters) and switches

\*\*\* Requires IOPless Fibre Channel adapters, IBM i 7.1 TR6 with resave RS-710-H released June 2013 and storage server level 6.4.1.4

@ BCH supports DS3400, DS3500, DS3950 & BCS supports DS3200, DS3500

% NPIV requires IBM i 7.1 TR2 (Technology Refresh 2) and latest firmware released May 2011 or later

%% NPIV requires IBM i 7.1 TR6 (Technology Refresh 6)

For a list of FC adapters supported by VIOS, Refer to the VIOS Datasheet at:  
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>.

Another source of information is the IBM i POWER External Storage Support Matrix Summary at  
<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4563>

To verify the supported combinations of adapters, servers and SANS use the System Storage Interoperation Center (SSIC) at:  
<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

### **3.4 Software and firmware: Direct attached**

<b>Software or firmware type (minimums)</b>	<b>Supported by IBM i for direct attachment</b>	<b>Notes</b>
IBM i 6.1 with LIC 6.1.1 or higher	Yes	See notes in section 3.3
IBM i 6.1 without LIC 6.1.1	No	
IBM i 5.4 or earlier	No	

### **3.5 Software and firmware: VIOS hosted**

<b>Software or firmware type (minimums)</b>	<b>Supported by IBM i as a client of VIOS</b>	<b>Notes</b>
IBM i 6.1 or higher	Yes	See notes in section 3.3
IBM i 5.4 or earlier	No	
PowerVM Standard Edition or higher	Yes	Required for VIOS media and activation code
VIOS 1.5 with Fix Pack 10.1 or higher	Yes	This is the minimum level of VIOS required
IBM Power server system firmware 320_040_031 or higher	Yes	This is the minimum system firmware level required
HMC firmware HMC V7 R3.2.0 or higher	Yes	This is the minimum HMC firmware level required
AIX/VIOS host kit license for DS3000/4000/5000	Yes	This license is required in order to attach DS3000/4000/5000 VSCSI LUNs to VIOS, and therefore virtualize them to IBM i.

## 4 IBM i using direct attach storage

The capability to use open storage directly attached to IBM i has been around for a while, but only on SANs that can support 520 byte sectors such as the DS8000's. In more recent releases of IBM i, the capability to handle the processing of 512 byte sectors has expanded the list of SANs that can be directly attached. The DS5300 and DS5100 are on this list with the introduction of IBM i 7.1 TR2. With additional PTFs on IBM i 7.1 TR6 shipped in 2Q 2013, we can also directly attach the Storwize products (SVC, V7000, V5000, V3700 and V3500) to IBM i. The following sections address a high level view of the configuration steps involved.

### 4.1 IBM i SAN configuration steps

#### 4.1.1 IBM i host kits

On the DS5100 and DS5300 you have to purchase and install an IBM i host kit in order to map volumes to an IBM i host. DS8000 does not have a host kit concept. Storwize and SVC products do not have host kits.

#### 4.1.2 Determine the FC adapter WWPNs

The fibre channel adapter ports have unique worldwide port names (WWPNs) assigned by the manufacturer. The WWPNs identify the host that owns the adapter. You need to retrieve the wwpn(s) of the adapters from IBM i or the HMC. With IBM i V7R1 and above, the wwpn and other details are now available in the WRKHDWRSC details of a storage adapter. You can use the below command on HMC to get the WWPN's and MAC address of the managed system.

```
lshwres -r io --rsubtype slotchildren -m <managed system name> --filter "lpar_names=< IBM i partition name here >" -F phys_loc,description,wwpn
```

Following are the requirements.

- The HMC is at the V7.3.2 level or higher.
- The managed server has POWER6 processors or newer.
- The managed server firmware (FSP) level is current/recent.
- The hardware discovery option is chosen when the managed server is powered on.

#### 4.1.3 Configure SAN switch zoning

Use either the command line interface (CLI) or the graphical user interface (GUI) for the SAN switch(es) to configure zoning to tie the IBM i host's world-wide port name (WWPN) of the fibre channel adapter ports on the adapter that is owned by the IBM i partition to the SAN's host port WWPNs. Use the System Service Tools, Hardware Service Manager interface to drill down to the FC adapter level to see the WWPNs.

#### 4.1.4 IBM i host configuration

Use either the command line interface (CLI) or the graphical user interface (GUI) for the SAN to configure a host connection tied to the WWPN. Doing the zoning first should help facilitate the host configuration.

#### 4.1.5 Configure SAN storage for IBM i

Both SANs support RAID sets across multiple disk drives. Choose the optimal RAID level for your IBM i workload. The result is a pool or array of storage that can be divided into logical volumes or LUNs. Even numbers of LUNs and similar sized LUNs perform well.

#### **4.1.6 Map the SAN LUNs to the IBM i host**

Once the LUNs are created, map them to the IBM i host that you created above. Use either the CLI or GUI.

#### **4.2 Add drives to the IBM i ASP**

Now the LUNs should be seen from IBM i. Use the System Service Tools interface to see the new drives and to add them to the targeted auxiliary storage pool (ASP).

## 5 IBM i hosted by VIOS

### 5.1 Virtual SCSI and Storage virtualization concepts

The IBM i storage portfolio now includes integrated SCSI, SAS or Solid State disk (SSD), as well as FC-attached storage subsystems that support direct attachment to IBM i, as discussed in the prior section. The capability to use open storage through VIOS extends the IBM i storage portfolio to include other 512-byte-per-sector storage subsystems. The logical units (LUNs) or volumes created on open storage are mapped through the FC fabric to VIOS. In VIOS the LUNs are virtualized to the IBM i virtual client partition. Figure 4 illustrates both the physical and virtualization components of the solution:

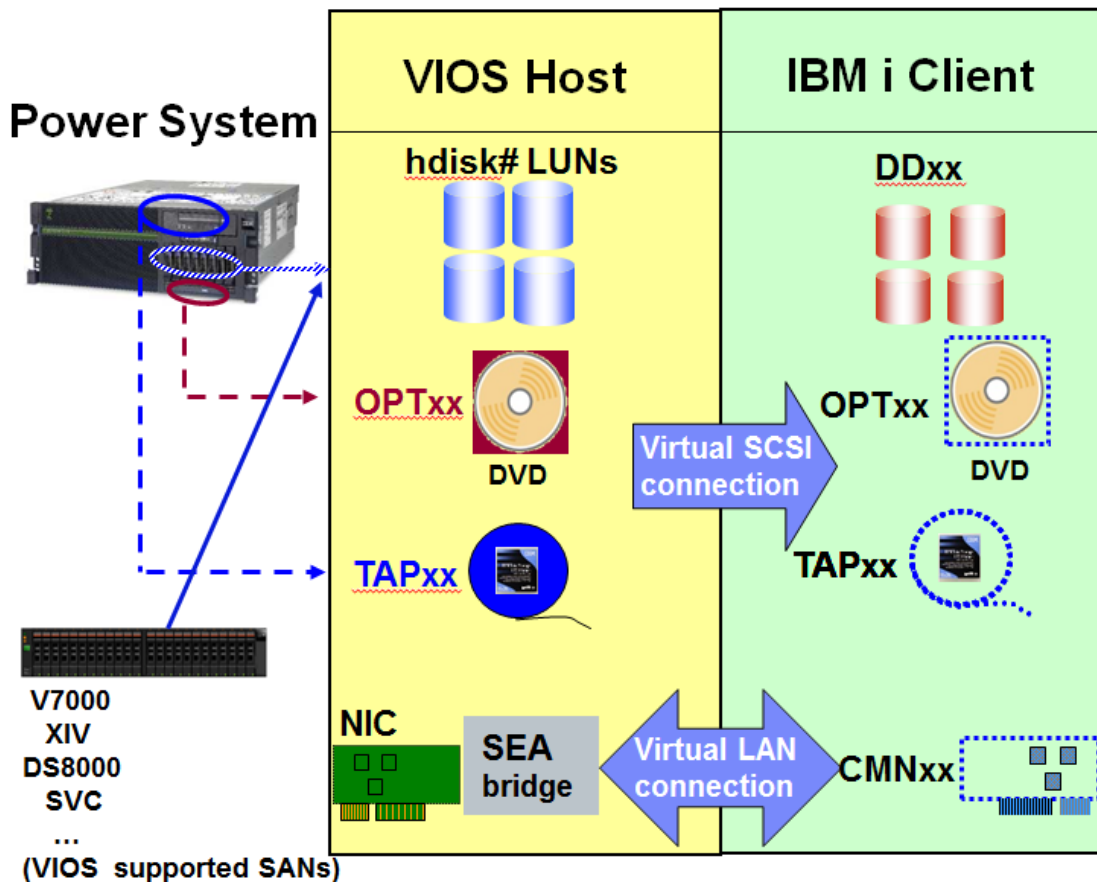


Figure 4: physical and virtualization components of the solution

VIOS has been used successfully for several years to virtualize storage, optical and networking resources to AIX and Linux client partitions. Now IBM i joins this virtualization environment, gaining the ability to use more types of open storage. From a VIOS perspective, IBM i is another client partition; the host-side configuration steps are the same as for AIX and Linux clients. While code changes were made in VIOS to accommodate IBM i client partitions, such as the disk sector conversion from 512 bytes to 520 bytes, there is no special version of VIOS in use for IBM i. If you have existing skills in attaching open storage to VIOS and virtualizing I/O resources to client partitions, they will continue to prove useful when creating a configuration for an IBM i client partition.

The hardware and virtualization components for attaching open storage to IBM i illustrated in Figure 4 also apply to using DS5000, XIV and other subsystems supported for this solution, as listed in the “IBM i using open storage configurations” section.

Three management interfaces are available for virtualizing I/O resources through VIOS to IBM i: the Hardware Management Console (HMC), the Systems Director Management Console (SDMC) and the Integrated Virtualization Manager (IVM). All provide logical partitioning and virtualization management functions, however, their support on Power servers and Power blades is different:

- The HMC is supported with all POWER6 and POWER7 servers and with the Nov 2012 firmware it supports Power blades.
- The SDMC is supported with all POWER6 and POWER7 servers and is supported with Power blades
- IVM is supported with all Power blades and with POWER6 and POWER7 express servers.

There is another significant difference between HMC and IVM when managing I/O virtualization for IBM i: when the Power server is IVM-managed, IBM i partitions must be purely virtual. They are not allowed to own any physical I/O adapters. Therefore, when managing Power servers, an HMC or SDMC is used most of the time, with IVM being used occasionally only on low end servers. SDMC is just being introduced as this paper is being updated. When managing Power blades, IVM or SDMC can be used. With the Nov 2012 combination of HMC firmware and FSP firmware, the HMC can manage Power blades too.

Note that using virtual I/O resources from VIOS does not preclude an IBM i client partition from owning physical hardware when the server is managed by an HMC or and SDMC, except for a POWER blade where it is not supported. A mix of virtual and physical hardware in the same partition is supported for IBM i in this environment, by assigning both types of adapters to the partition in the HMC/SDMC.

## ***5.2 Virtual Fibre Channel adapter concepts***

NPIV is an industry-standard FC protocol that allows VIOS to directly share a single FC adapter among multiple client LPARs. NPIV can be used for disk storage and or tape attachment. Unlike VSCSI, NPIV does not map a LUN or tape device to a virtual target device in VIOS, which the client LPAR can then access as a generic SCSI disk or a tape device. Instead, a port on the physical FC adapter is mapped to a virtual FC server adapter in VIOS, which in turn is connected to a virtual FC client adapter in IBM i.

When the virtual FC client adapter is created, two unique world-wide port names (WWPNs) are generated for it. Through the link to the server virtual FC adapter and then the physical adapter in VIOS, those WWPNs become available on the SAN and storage and or tape drives can be mapped to them as with any other FC host ports. Note that these WWPNs are unique not just within the Power server, but globally on the SAN. The PowerVM Hypervisor creates two WWPNs per virtual FC client adapter to facilitate Live Partition Mobility for LPARs using NPIV. When a virtual FC client adapter is deleted, the WWPNs are not reused. By default, the PowerVM Hypervisor is capable of creating 32,000 WWPNs for virtual FC client adapters. If additional WWPNs are required, clients can acquire an enablement code from IBM.

Figure 5 illustrates using NPIV for IBM i disks:



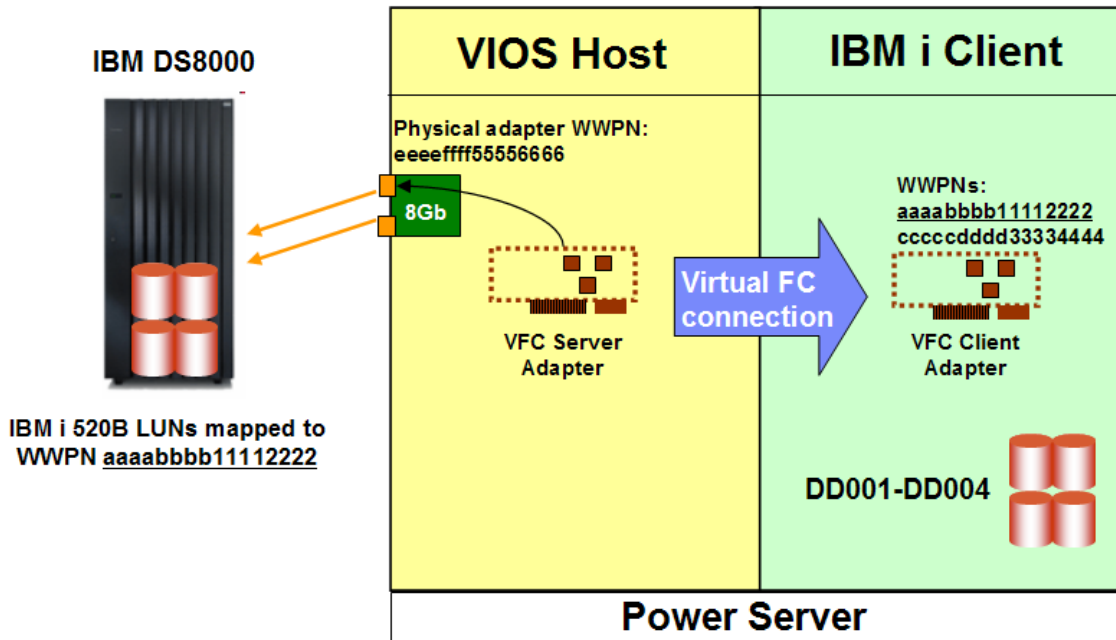


Figure 5: Using NPIV for IBM i

From the storage subsystem's perspective, the LUNs, volume group and host connection are created as though the IBM i LPAR is directly connected to the storage through the SAN fabric. While VIOS still plays a role with NPIV, that role is much more of a passthrough one when compared to VSCSI. Note that an 8Gb or 16Gb FC adapter is required; however, the FC switches do not need to be at that speed. Additionally, only the first switch – the one to which the FC adapter on the Power server is directly connected – must be NPIV-capable. The rest of the SAN fabric does not need to be NPIV-capable.

### 5.2.1 NPIV Supported hardware and software

Use the following link to determine the supported hardware and required software for NPIV attached tape devices for IBM i: <https://www.ibm.com/developerworks/ibmi/media>

### 5.2.2 Creating Virtual Fibre Channel (VFC) adapters using the HMC

When using an HMC to manage Power Servers or Blades, you can create fibre channel adapters to create the virtual fibre channel connection between VIOS and the IBM i client partition. In this section you will see how to create a server adapter and client adapter pair. You need to determine the next adapter slot that is available for both the VIOS and the IBM i partitions. The slot numbers are used to tie the adapters together.

Use the **Current configuration** for the VIOS and IBM i partition's to **note** the next available virtual adapter slot number. This will be used in some of the commands that follow.

Create Server Virtual Fibre Channel Adapter

Perform the following steps to create a server fibre channel adapter.

- Log into HMC with **hscroot** as the user ID or using another user ID.
- Click the **Managed System** link of the navigation pane
- Select the server from the list and select the checkbox next to the VIOS you are configuring
- Click **Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit...** to edit the profile

- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> Fibre Channel Adapter...**
- Specify an available **Server Adapter slot number** and click “**This adapter is required for partition activation**” if this adapter will be used for disk that will be needed to boot the IBM i partition.
- Use the list to select the **IBM i client partition** you want to use
- Specify an available **Client Adapter ID** (slot number) and click **OK** to create the server Virtual Fibre Channel Adapter

Note: An alternative is to create the Server Virtual Fibre Channel adapter dynamically if VIOS is currently running. In this case, replace the above steps with the following:

- Select the server from the list and select the checkbox next to the VIOS you are configuring
- Click **Dynamic partitioning -> Virtual Adapters**
- **Select Actions in the upper left corner of the window and then click Actions -> Create Virtual Adapter -> Fibre Channel Adapter...**
- Specify an available **Server Adapter slot number**.
- Use the list to select the **IBM i client partition** you want to use
- Specify an available **Client Adapter ID** and click **OK**.
- **Then click OK one more time** to create the server Virtual Fibre Channel Adapter dynamically

Create Client Virtual Fibre Channel Adapter

Perform the following steps to create a client fibre channel adapter:

- Select the server from the HMC and click the checkbox next to the IBM i client you are configuring.
- Click **Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit...** to edit the profile
- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> Fibre Channel Adapter...**
- Specify an available **Client Adapter slot number** and select the **VIOS** from the Server partition list.
- Specify an available **Server Adapter ID** (slot number) and click **OK** to create the client Virtual Fibre Channel Adapter.

Note: the adapter slot numbers in the client profile need to match the adapter slot numbers specified in the VIOS profile.

### 5.2.3 Mapping VFC adapters to physical ports using the HMC

- Sign onto the HMC as hscroot or another superadministrator-level user ID.
- Expand **Systems Management**
- Click **Servers**
- Select the correct managed server
- In the menu below, expand **Configuration** and then **Virtual Resources**
- Click **Virtual Storage Management**
- Select the correct VIOS from the list and click **Query VIOS**
- Select the **Virtual Fibre Channel** tab.
- Select a physical port fcsN and click **Modify partition connections...**
- Check the box next to the IBM i partition that you want to map storage to and click **Ok**

### 5.3 Virtual SCSI (VSCSI) configuration concepts

For VIOS to virtualize LUNs created on open storage to an IBM i client partition, configuration takes place both in the HMC, SDMC or IVM and in VIOS. No explicit VSCSI configuration is necessary in IVM the majority of the time.

In the HMC/SDMC, the minimum required configuration is:

- One VSCSI server adapter in the host partition
- One VSCSI client adapter in the client partition

This VSCSI adapter pair allows the client partition to send read and write I/O operations to the host partition. More than one VSCSI pair can exist for the same client partition in this environment. To minimize performance overhead in VIOS, the VSCSI connection is used to send I/O requests, but not for the actual transfer of data. Using the capability of the POWER Hypervisor for Logical Remote Direct Memory Access (LRDMA), data is transferred directly from the FC adapter in VIOS to a buffer in memory of the IBM i client partition.

In an IBM i client partition, a VSCSI client adapter is recognized as a type **290A DCxx** storage controller device. Figure 6 depicts the VSCSI client adapter, as well as several open storage LUNs and an optical drive virtualized by VIOS:

```
Work with Storage Controller Resources
System: Y213B2P2
Type options, press Enter.
5=Work with configuration descriptions 7=Display resource detail

Opt Resource      Type-model      Status          Text
-  DC01            290A-001        Operational     Storage Controller
-  DD001           6B22-050        Operational
-  DD002           6B22-050        Operational
-  DD003           6B22-050        Operational
-  DD004           6B22-050        Operational
-  DD005           6B22-050        Operational
-  OPT01          632C-002        Operational     Optical Storage Unit
-  DD006           6B22-050        Operational
-  DD007           6B22-050        Operational
```

Figure 6: VSCSI client adapter and open storage LUNs and an optical drive as seen in IBM i.

In VIOS, a VSCSI server adapter is recognized as a **vhostX** device:

```
$ lsdev | grep vhost
vhost0          Available      Virtual SCSI Server Adapter
$
```

Figure 7: Example of VSCSI adapter seen as vhost0 by VIOS.

In VIOS, however, a new object must be created for each open storage LUN that will be virtualized to IBM i: a *virtual target SCSI device*, or *vtscsiX*. A vtscsiX device makes a storage object in VIOS available to IBM i as a standard DDxxx disk unit. There are three types of VIOS storage objects that can be virtualized to IBM i:

- Physical disk units or volumes (*hdiskX*), which are open storage LUNs in this case
- Logical volumes (*hdX and other*)
- Files in a directory

For both simplicity and performance reasons, **it is recommended to virtualize open storage LUNs to IBM i directly as physical devices (hdiskX)**, and not through the use of logical volumes through VIOS. (Refer to the "" section for a detailed performance discussion.) A vtscsiX device links a LUN available in VIOS (hdiskX) to a specific VSCSI adapter (vhostX). In turn, the VSCSI adapter in VIOS is already connected to a client SCSI adapter in the IBM i client partition. Thus, the hdiskX LUN is made available to IBM i through a vtscsiX device. What IBM i storage management recognizes as a DDxxx disk unit is not the open storage LUN itself, but the corresponding vtscsiX device. The vtscsiX device correctly reports the parameters of the LUN, such as size, to the virtual storage code in IBM i, which in turn passes them on to storage management.

Multiple vtscsiX devices, corresponding to multiple open storage LUNs, can be linked to a single vhostX VSCSI server adapter and made available to IBM i. It is possible to virtualize up to 16 LUNs to IBM i through a single VSCSI connection. This is an IBM i limit, not VIOS. NOTE: As of i7.1 TR8 this limit has been increased to 32. Each LUN typically uses multiple physical disk arms in the open storage subsystem. If more than 16/32 LUNs are required in an IBM i client partition, an additional pair of VSCSI server (VIOS) and client (IBM i) adapters must be created in the HMC/SDMC or on the IVM/VIOS command line. Additional LUNs available in VIOS can then be linked to the new vhostX device through vtscsiX devices, making them available to IBM i.

Prior to May 2009, creating vtscsiX devices and thus virtualizing open storage LUNs to IBM i was necessarily a task performed only on the IVM/VIOS command line when the server was HMC-managed. When the server is IVM-managed, assignment of virtual resources is performed using the IVM browser-based interface for the first 16/32 devices for a client partition. However, the HMC interface has been significantly enhanced, including support for assigning vtscsi devices to multiple VSCSI adapters. This support became available on HMC firmware Version V7.7.3.x.

For HMC/SDMC provisioning of VIOS storage to work, the VIOS LPAR must have an active Resource Monitoring and Control (RMC) connection to the HMC. The same connection is required for performing dynamic LPAR (DLPAR) changes to the VIOS LPAR and it generally involves successful TCP/IP communication between the LPAR and the HMC. If the HMC displays an RMC error while attempting VIOS storage provisioning, refer to this document: <http://www.ibm.com/support/docview.wss?uid=isg3T1012915>

If the storage configuration is performed while VIOS is down or not yet installed, new storage will be discovered upon boot and the HMC can then be used to assign the correct LUNs to IBM i. If storage changes are made while VIOS and IBM i are running, it will still be necessary to use Telnet or PuTTY to VIOS and run the **cfgdev** command to discover newly added storage prior to using the HMC for assigning it to IBM i.

### 5.3.1 VSCSI adapter configuration using IVM

VIOS providing storage to an IBM i client LPAR uses an existing function of the system firmware, or POWER Hypervisor: the capability to create VSCSI and Ethernet adapters in a partition. Virtual adapters are created for each LPAR in the HMC/SDMC. VSCSI adapters are used for storage and optical virtualization.

Using IVM for LPAR and virtualization management mandates using only virtual resources for IBM i partitions. However, IVM provides the benefit of automatically creating VSCSI server (in VIOS) and client adapters (in IBM i) as virtual disk, optical and tape resources are assigned to IBM i. For more information on LPAR management with IVM, refer to the IVM topic in the *PowerVM Editions Guide*, available at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm>.

Note that IVM will automatically create only the first pair of VSCSI server and client adapters when disk and optical resources are virtualized to IBM i. Additional adapter pairs must be created by using the IVM command line, available through a Telnet or Secure Shell (SSH) session to VIOS. Refer to the VIOS topic in the *PowerVM Editions Guide* for more information on using the IVM/VIOS command line. IVM will automatically create a new VSCSI adapter pair for each tape resource assigned to IBM i.

HMC and SDMC offer interfaces to manually create virtual adapters and does not create them by default as IVM does. Please see the following sections for more details.

As discussed above, IVM creates a pair of VSCSI server and client adapters as you map disks, tapes or optical devices through the browser interface to an IBM i partition. Additional adapter pairs have to be created using the VIOS command line interface (CLI).

### 5.3.2 Creating multiple Virtual SCSI adapters using the VIOS CLI

To create additional Virtual SCSI client adapters under IVM, you must use the VIOS command line:

- Log into VIOS with **padmin** or another administrator user ID.

VIOS always has partition ID 1 when IVM is used, and by default carries the serial number of the blade as a name.

To display the current names and IDs of all existing partitions, use the following command:

- **lssyscfg -r lpar -F "name,lpar\_id"**

To display which virtual slots have been already been used, by partition and adapter, use:

- **lshwres -r virtualio --rsubtype slot --level slot** (notice the double dashes)
- **The type of adapter is shown on the middle of each line. The slot numbers are shown at the beginning of each line.**

If the IBM i partition is not activated, refer to the following example that adds a new Virtual SCSI client adapter in **slot 5** of IBM i partition **"test,"** connecting to a server adapter in the next available slot (chosen automatically by IVM) in the partition named **"VIOS"**:

- **chsyscfg -r prof -i "name=test,virtual\_scsi\_adapters+=5/client/1/VIOS/1"**
- The corresponding server adapter in VIOS is created automatically by IVM
- If the IBM i partition is running, refer to the following example that creates a new Virtual SCSI client adapter in **slot 5** of IBM i partition **"test,"** connecting to a server adapter in the next available slot (chosen automatically by IVM) in partition **"VIOS"**:
  - **chhwres -r virtualio --rsubtype scsi -p test -o a -s 5 -a "adapter\_type=client"**
  - The corresponding server adapter in VIOS is created automatically by IVM
- Notice that there are three variables in the previous commands: the name of the IBM i partition, the new slot for the Virtual SCSI client adapter, and the name of the VIOS partition.

### 5.3.3 Creating multiple Virtual SCSI adapters using the HMC

When using an HMC to manage Power Servers or Blades, you can create virtual SCSI adapters to create the virtual SCSI connection between VIOS and the IBM i client partition. In this section you will see how to create server adapter and client adapter pair. You need to determine the next adapter slot that is available for both the VIOS and the IBM i partitions. The slot numbers are used to tie the adapters together.

Use the **Current configuration** for the VIOS and IBM i partition's to **note** the next available virtual adapter slot number. This will be used in some of the commands that follow.

#### Create Server SCSI Adapter

Perform the following steps to create a server SCSI adapter.

- Log into HMC with **hscroot** as the user ID or using another user ID.
- Click the **Managed System** link of the navigation pane
- Select the server from the list of hosts and select the checkbox next to the VIOS you are configuring
- Click **Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit...** to edit the profile
- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> SCSI Adapter...**
- Specify the next available **Server Adapter slot number** (determined above) and click "**Only selected client partition can connect**". Do NOT use Any Partition can connect. That will not work.
- Use the list to select the **IBM i client partition** you want to use
- Specify the next available (determined above) **Client Adapter slot number** and click **OK** to create the server Virtual SCSI Adapter

#### Create Client SCSI Adapter

Perform the following steps to create a client SCSI adapter:

- Select the server from the HMC and click the checkbox next to the IBM i client you are configuring.
- Click **Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit...** to edit the profile
- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> SCSI Adapter...**
- Specify the next available **Client Adapter slot number** (determined above) and select the **VIOS** from the Server partition list.
- Specify the next available **Server Adapter slot number** (determined above) and click **OK** to create the client Virtual SCSI Adapter.

### 5.3.4 Creating multiple Virtual SCSI adapters using the SDMC

When using an SDMC to manage Power Blades, you can create virtual SCSI adapters to create the virtual SCSI connection between VIOS and the IBM i client partition. In this section you will see how to create server adapter and client adapter pair. You need to determine the next adapter slot that is available for both the VIOS and the IBM i partitions. The slot numbers are used to tie the adapters together.

Use the **Current configuration** for the VIOS and IBM i partition's to **note** the next available virtual adapter slot number.

#### Create Server SCSI Adapter

Perform the following steps to create a server SCSI adapter.

- Log into SDMC with **sysadmin** as the user ID or using another user ID.
- From the **Resources** tab of the SDMC, select the blade from the list of hosts and select the checkbox next to the VIOS you are configuring.
- Click **Actions-> System Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit...** to edit the profile

- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> SCSI Adapter...**
- Specify the next available (determined above) **Server Adapter slot number** click “**Only selected client partition can connect**”
- Use the list to select the IBM i client partition you want to use
- Specify the next available (determined above) **Client Adapter slot number** and click **OK** to create the server Virtual SCSI Adapter

Create Client SCSI Adapter

Perform the following steps to create a client SCSI adapter:

- Click the Resources tab of the SDMC welcome screen and select the checkbox next to the IBM i client you are configuring
- Click **Actions-> System Configuration -> Manage Profiles**
- Select a profile and select **Actions->Edit...** to edit the profile
- Click the **Virtual Adapters** tab and then click **Actions -> Create Virtual Adapter -> SCSI Adapter...**
- Specify the next available (determined above) **Client Adapter slot number** and select the VIOS from the Server partition list.
- Specify the next available (determined above) **Server Adapter slot number** and click **OK** to create the client Virtual SCSI Adapter

As soon as they are available in VIOS, open storage LUNs and optical devices can be assigned to IBM i using one of the following options:

### 5.3.5 Assigning storage using IVM

- Click on **View/Modify Virtual Storage**.
- Click the **Physical Volumes** tab.
- Select the hdisks that you want to assign to the IBM i client.
- Click **Modify Partition Assignment**.
- Select the IBM client partition from the list and click **OK**.

### 5.3.6 Assigning storage using the HMC

- Sign onto the HMC as hscroot or another superadministrator-level user ID.
- Expand **Systems Management**.
- Click **Servers**.
- Select the correct managed server.
- In the menu below, expand **Configuration** and then **Virtual Resources**.
- Click **Virtual Storage Management**.
- Select the correct VIOS from the list and click **Query VIOS**.
- Click the **Physical Volumes**.
- Select the correct LUNs (up to 16, i7.1 TR8 has increased this value to 32) or optical devices (up to 16/32) and click **Modify assignment...** to assign them to IBM i. Additional LUNs can be selected, as the AIX limit is 256, but IBM i will not be able to see them.
- **Note:** For more than 16/32 LUNs, look for the pull down in the upper middle of the **Modify assignment...** screen. If you have created additional VSCSI adapters for your client partition, they will be listed in the pull down. So you can select another adapter and map drives to it. You can also still use the make virtual device (mkvdev) command from the VIOS command line interface if you prefer that interface.

Note that this sequence assumes the existence of the correct VSCSI server and client adapters in VIOS and IBM i, respectively. After the assignment is performed, the new LUNs are immediately available in IBM i as non-configured drives.

### 5.3.7 Assigning storage using SDMC

SDMC can also be used to assign hdisks to the **IBM i client** partitions (up to 16 per VSCSI adapter pair (i7.1 TR8 this limit has been increased to 32). Perform the following steps to assign storage using SDMC.

- On the Resources link of the welcome page, expand **Hosts** and click on the host (Power server) that you want to work with.
- Select the virtual server for VIOS, right-click on the name and click **System Configuration -> Manage Virtual Server**.
- Click **Storage Devices** on the left-hand navigation panel. The physical volumes (hdiskX) are shown on the right.
- Select the hdisks that you want to assign to the IBM i client and click **Assign**.
- Select the client IBM i partition from the list and click **OK**.

One use of the command line is to add more details to the LUNs to better manage them. The make virtual device (mkvdev) command has a device parameter that can be used to add more descriptive names. Refer to the VIOS command reference for more details at:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/iphcg/iphcg.pdf>

If the QAUTOCFG system value in IBM i is set to 1 (which is the default), the new virtual resources will become available in IBM i immediately. No action is required on the IVM/VIOS command line.

## 5.4 Optical virtualization

### 5.4.1 Physical optical devices

CD and DVD drives supported by VIOS can be virtualized to IBM i directly. Optical drives are recognized in VIOS as cdX devices. A cdX device is linked to a vhostX VSCSI server adapter and made available to IBM i through a *virtual target optical device*, *vtoptX*. The same optical drive can be virtualized to multiple IBM i client partitions by creating separate vtoptX devices linked to different vhostX adapters for the same cdX device. Only one IBM i client partition can use the physical optical drive at a time. If the physical optical drive is writeable, IBM i will be able to write to it. Similar to LUNs, optical devices can be virtualized to IBM i using the enhanced functions of the HMC/SDMC; it is not necessary to use the IVM/VIOS command line to create the vtoptX devices.

### 5.4.2 VIOS Media Repository

VIOS provides a capability similar to that of an image catalog (IMGCLG) in IBM i: a repository of media images on disk. Unlike IMGCLGs in IBM i, a single media repository may exist per VIOS. The media repository allows file-backed virtual optical volumes to be made available to the IBM i client partition through a separate vtoptX device. One immediate benefit to IBM i is the ability to import the IBM i install media into VIOS as ISO images, then install the client partition from the ISO images instead of switching real media in a physical DVD drive.

### 5.4.3 Assigning a physical or virtual Optical/DVD/CD drive using the HMC

The HMC also has an interface to create the virtual media library:

- You need to start at the Managed Server level. Select the managed server and then expand **Configuration->Virtual Resources->Virtual Storage Management**.
- Select a VIOS server and click **Query VIOS**.



- Expand the **Optical devices** tab and select the **cd#/virtual optical device** and assign it to a partition.

#### 5.4.4 Assigning a physical or virtual Optical/DVD/CD drive using the SDMC

SDMC has an interface to create and load the virtual media library too. Use the following steps to do this:

- On the welcome page, click **Resources**,
- Expand **Hosts** and click on the host (Power server) that you want to work with.
- Select the virtual server for VIOS, right-click the name and click **System Configuration - > Manage Virtual Server**.
- Click **Media Devices** on the left-hand navigation panel. The optical device interface is shown on the right.
- Select the **cd#/virtual optical device** and assign it to a partition.

#### 5.4.5 Assigning a physical or virtual Optical/DVD/CD drive using IVM

IVM has an interface to move the optical drive under the View/Modify Storage task, then clicking the Optical/CD tab. Select the **cd#/virtual optical device** and assign it to a partition.

### 5.5 Network virtualization

Using virtual Ethernet adapters for a partition-to-partition communication within a Power server is an existing Power server capability. In order for an IBM i client to use a physical network adapter in VIOS, a virtual Ethernet (VE) adapter must be created in both of the partitions in the HMC/SDMC. The virtual Ethernet adapters must have the same Port Virtual LAN ID (PVID) in order to communicate. The IBM i VE is typically configured as part of the partition configuration wizard. But there may be a need to later add another VE adapter. See the following sections for details on how to do this.

#### 5.5.1 Configuring Virtual Ethernet using the HMC

Perform the following steps in the HMC to add a virtual Ethernet adapter to a partition:

- Select the managed server and left click on its name.
- Select the partition you want to change by clicking the selection box next to its name.
- In the menu below, expand **Dynamic partitioning**, then **Virtual adapters**.
- Using the **Actions** pulldown in the upper left hand corner, select **Create virtual adapter - > Virtual Ethernet**.
- The Adapter ID is filled in automatically. Enter the **Port Virtual Ethernet (VLAN ID)** that you want the partition to use. Each partition that wants to talk on this VLAN should use the same VLAN ID.
- IBM i does not support VLAN tagging, but does allow the IEEE 802.1q compatible adapter type to be selected to separate traffic from VIOS.
- If this VE adapter is for VIOS and it will be used for a shared Ethernet adapter (SEA) (coming up in the following sections), then select the **Use this adapter for Ethernet bridging**.
- Click **OK** to create the adapter.
- This was a dynamic add to the partition. You should go under **Configuration -> Save Current Configuration** to save the changes to the partition profile.

## 5.5.2 Configuring Virtual Ethernet using the SDMC

The Create Virtual Server wizard queries the available Ethernet connections of the VIOS. These are shown during the VE configuration step. You can choose the adapter or adapters to add to your IBM i virtual server. If you need to add another VE adapter, perform the following steps to do this:

- On the welcome page click **Resources**,
- Expand **Hosts** and click on the host (Power server) that you want to work with.
- Select the virtual server. **Right-click** on the name and click **System Configuration -> Manage Virtual Server**.
- Click **Network** on the left-hand navigation panel. The virtual network interface is shown on the right.
- The Adapter ID is filled in automatically. Enter the **Port Virtual Ethernet (VLAN ID)** that you want the partition to use. Each partition that wants to talk on this VLAN should use the same VLAN ID.
- IBM i does not support VLAN tagging, but does allow the IEEE 802.1q compatible adapter type to be selected to separate traffic from VIOS.
- If this VE adapter is for VIOS and it will be used for a shared Ethernet adapter (SEA) (coming up in the following sections), then select the **Use this adapter for Ethernet bridging**.
- Click **OK** to create the adapter.
- This was a dynamic add to the partition. You should go under **Configuration -> Save Current Configuration** to save the changes to the partition profile.

Additionally, link aggregation is supported for fail over and load balancing across multiple Ethernet ports. The **cfgassist** menu interface can be used from the VIOS CLI to configure link aggregation, SEAs and TCP/IP.

## 5.5.3 Configuring Virtual Ethernet using IVM

To configure the Virtual Ethernet adapter, start an IVM session and log in with **padmin** user ID.

- Click **View/Modify Virtual Ethernet**.
- Select the **Virtual Ethernet LAN** tab
- Notice the Virtual Ethernets that are supported by VIOS, shown by a checkmark.
- Select one of those same virtual Ethernets for your IBM i client partition.
- Click **Apply**

## 5.5.4 Configuring Virtual Ethernet using the VIOS CLI

To create additional Virtual Ethernet adapters you must use the VIOS command line:

- Log into VIOS with **padmin** or another administrator user ID.

VIOS always has partition ID 1 when IVM is used, and by default carries the serial number of the blade as a name.

To display the current names and IDs of all existing partitions, use the following command:

- **Issyscfg -r lpar -F "name,lpar\_id"**

To display which virtual slots have been already been used, by partition and adapter, use:

- **Ishwres -r virtualio --subtype slot --level slot** (notice the double dashes)
- **The type of adapter is shown on the middle of each line. The slot numbers are shown at the beginning of each line.**

If the IBM i partition is not activated, refer to the following example that adds a new Virtual Ethernet adapter to IBM i partition “test,” in **slot 5**, is not IEEE 802.1Q enabled (it does not support VLAN tagging), has a port virtual LAN ID (PVID) of 50, no additional virtual LAN IDs, it is not a trunk adapter, and is required.

- **chsyscfg -r prof -i "name=test,virtual\_eth\_adapters=5/0/50///1"**

If the IBM i partition is running, refer to the following example to add a virtual Ethernet adapter to partition 4 in virtual slot 5 with VLAN tags on VLAN 212 and VLAN 313:

- **chhwres -r virtualio --subtype eth -o a --id 4 -s 5 -a port\_vlan\_id=2,ieee\_virtual\_eth=1,\"addl\_vlan\_ids=212,313\"**

For other examples of configuring VE adapters with tagged VLANs see:

<http://www.ibm.com/support/docview.wss?uid=isg3T1010975> or  
<http://www.ibm.com/support/docview.wss?uid=isg3T1012722>

### 5.5.5 Shared Ethernet Adapters (SEAs)

In VIOS, the same Ethernet type of device, entX, is used for logical Host Ethernet ports, physical and virtual Ethernet adapters:

```
$ lsdev | grep ent
ent0          Available Logical Host Ethernet Port (lp-hea)
ent1          Available Logical Host Ethernet Port (lp-hea)
ent2          Available Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent3          Available Gigabit Ethernet-SX PCI-X Adapter (14106703)
ent4          Available Virtual I/O Ethernet Adapter (1-lan)
```

Figure 8: Example list of Ethernet ports in VIOS.

Inter-partition communication is facilitated via virtual Ethernet adapters using the same PVIDs. VIOS provides access to external networking to client partitions by bridging a physical Ethernet adapter and one or more virtual Ethernet adapters. The virtualization object that provides this Ethernet bridge is called a *Shared Ethernet Adapter (SEA)*. The SEA forwards network packets from any client partitions on a VLAN to the physical LAN through the physical Ethernet adapter. Because the SEA creates a Layer-2 bridge, the original MAC address of the virtual Ethernet adapter in IBM i is used on the physical LAN. The CMNxx communications port that represents the virtual Ethernet adapter in IBM i is configured with an externally routable IP address and a standard network configuration is used. The physical adapter bridge by the SEA can be any network adapter supported by VIOS, including Integrated Virtual Ethernet (IVE) ports, also known as Host Ethernet Adapter (HEA) ports.

NOTE: When configuring redundant VIOS's, another virtual Ethernet adapter must be created in each VIOS which will be used as the control channel. These virtual adapters must use a PVID that is different than the one used by the SEA and cannot be configured for Layer-2 bridging. See the section “Redundant VIOS Networking” for more details.

#### 5.5.5.1 Configuring an SEA using IVM

To configure the Shared Ethernet adapter (SEA), start an IVM session and log in with **padmin** user ID.

- Click **View/Modify Host Ethernet Adapter** or the applicable Ethernet hardware option
- Ensure the option for **Allow virtual Ethernet Bridging** is set.
- Click **View/Modify Virtual Ethernet**.
- Click the **Virtual Ethernet Bridge** tab.
- On the **Virtual Ethernet ID 1** select the physical adapter to the Ethernet port you want to bridge to.

- Click **Apply**

### 5.5.5.2 Configuring SEAs using the HMC/SDMC

Follow the steps listed in the above HMC/SDMC sections to get to the **virtual network** interface.

- Select the correct managed server
- In the menu below, expand **Configuration**, then **Virtual Resources**
- Click **Virtual Network Management**
  - Select the **VLAN ID** that is configured for the IBM i client partition the VIOS that will bridge it. When you do this, you are shown the participants on the VLAN.
  - Near the bottom right hand side of the window, select the **Create SEA...** button.
    - Choose the **VIOS(s)** where the SEA will be created.
    - Choose the **physical adapter** to bridge the IO to.
    - For redundant VIOS's, check the **Configure virtual network failover** box and select both the **Failover VIOS** and the **Failover physical adapter** values using the associated pull-downs.
    - Click **OK** to create the SEA.

### 5.5.5.3 Configuring SEAs using the VIOS CLI

The following document describes how to set up auto mode using the command line interface, but you can substitute sharing for the `ha_mode` parameter value. Refer to the following steps at: <https://www-304.ibm.com/support/docview.wss?uid=isg3T1011040>

Note that step 3 and step 7 in the above document are the `mkvdev -sea ...` command to create the shared Ethernet adapters (SEA) has to have all of the parameters, specifically the `-attr` parameters, included in the one command. Do not try to change an existing SEA to add the attributes.

Note: Step 5 Create a virtual adapter to be used in the SEA adapter on VIOS2. EX: (ent2)

- a. Give the virtual adapter the same VLAN ID (PVID) as VIOS1. EX: "1" .
- b. Check the box "access external network".
- c. Give the virtual adapter a higher trunk priority. EX: "2"

**The Trunk Priority of 2 is critical!** The default is 1, which will match the SEA already created in the first VIOS partition, with the same trunk priority, or if you try to create the second SEA without all of the attributes shown in step 7 of the referenced document, your VIO server will create a **broadcast storm** on your network! Your network administrator will not be happy with you.

Note: To better utilize the 10Gb Ethernet adapters that can be owned by each VIOS server, `ha_mode=sharing` was implemented. The virtual Ethernet traffic handling is negotiated between the VIOS servers as additional LANs are configured, but there is a catch. You have to enable IEEE 802.1Q VLAN tagging on each virtual Ethernet LAN for this to work. A different port VLAN ID (PVID) is not enough. You don't have to specify an explicit VLAN tag, just enable the tagging support on each virtual Ethernet adapter. Check the box labeled IEEE 802.1Q to enable tagging and have the negotiation take place.

## 5.6 Redundant VIOS Networking

When implementing redundant VIOS partitions as discussed in the Dual/Redundant hosting section, you also need to consider the networking used there. The default handling of the Ethernet traffic is for the primary VIOS to do all of it. This is referred to as `ha_mode=auto`. But users also wanted a way to share the handling of the Ethernet traffic between the primary and the

backup VIOS's. This is configured using `ha_mode=sharing`. You can use the HMC, SDMC or VIOS CLI do configure this.

See the steps in section “Configuring SEAs using the HMC/SDMC” to configure the SEA, but while on that interface look for a **Failover** checkbox.

- Select the **Failover** checkbox
- Then choose the **second VIOS partition** and its associated Ethernet adapter (**ent#**).
- Ensure that a **control channel**, used for the heartbeat function between the VIOS's, is selected to be configured. NOTE: When configuring a failover VIOS, another virtual Ethernet adapter must be created in each VIOS first which will be used as the control channel. These virtual adapters must use a PVID that is different than the one used by the SEA and cannot be configured for Layer-2 bridging.

## ***5.7 Adding new devices to VIOS***

VIOS does not automatically scan for new devices added after the initial boot is complete. When a new device – such as a LUN on a storage subsystem – is added to VIOS, prompt VIOS to scan for new devices by connecting with Telnet or SSH and using the following command:

- **cfgdev**
- If you are using IVM, click **Hardware Inventory** and then click **Configure devices** to run the `cfgdev` command.

After detecting the device, VIOS will automatically configure it and make it available for use. It can then be virtualized to the IBM i client partition using the HMC, as described in the previous section. Note that if the SAN configuration is performed before the VIOS partition boots, this step is not necessary, as VIOS will recognize all available devices at boot time.

## 6 Attaching open storage to IBM i through VIOS: Best Practices

### 6.1 Storage planning

The first storage planning consideration is to have sufficient available capacity in the open storage subsystem to create the AIX/VIOS LUNs that will be virtualized to IBM i by VIOS. As mentioned in the supported hardware the “IBM i using open storage configurations section, it is **strongly recommended that only FC or SAS physical drives are used to create LUNs for IBM i as a client of VIOS**. The reason is the performance and reliability requirements of IBM i production workloads. For non-I/O-intensive workloads or near line storage, SATA or FATA drives can also be used. This recommendation is not meant to preclude the use of SATA or FATA drives for other clients of VIOS or other host servers; it applies only to production IBM i workloads.

The storage subsystem may require a host kit for AIX/VIOS to be installed before attaching LUNs to these hosts (as is the case with DS5000 or DS4000) for mapping the drives to IBM i using the VSCSI adapters.

As the LUNs are virtualized by VIOS, they do not have to match IBM i integrated disk sizes. The technical minimum for any disk unit in IBM i is 160 MB and the maximum is 2 TB, as measured in VIOS. Actual LUN size is based on the capacity and performance requirements of each IBM i virtual client partition and load source disk restrictions.

A *path* to a LUN is a logical connection from the host HBA to the external storage box. The path may include virtual or physical HBAs and may have no switches or multiple switches. IBM i enforces up to 8 active *paths* to the *same* LUN. This is due to the overhead of allocating configuration resources for each path to the IBM i. Here are some ways that paths are created:

1. A direct cable from the host HBA to the external storage HBA will be 1 path.
2. Cabling a host FC adapter ports to a FC switch and zoning any of those adapter ports to a SAN host port will generate a path per world wide node name (WWNN). The IBM i host port will only log into one target port per WWNN even if there are multiple target ports cabled and zoned. The DS8000 is one WWNN where as the SVC and Storwize have 2 WWNNs per controller enclosure. All paths to the DS8000 are considered active where as the SVC or Storwize has one preferred controller that is considered the active path and the other controller will be the passive path. This includes native attached ports or VIOS (NPIV) attached ports.
3. Cabling multiple host ports to a switch and zoning those host ports to the IBM i will generate paths for every host port zoned to the partition. This includes native attached ports or VIOS (NPIV) attached ports.

Remember that the physical path must still be sized to provide adequate throughput. Since you have an 8 or 16 Gb physical port for the NPIV adapter you will have performance problems if too many clients attempt to use the NPIV adapter at the same time.

Note: Each adapter can support up to 128 total paths, but only 64 active paths. This limits the number of devices to 64. The only case where you should have more than 64 paths is when you use SVC or Storwize where you can also have the inactive path configured on the same host port as the active path. The 64 device/active path limit also applies to tape where the control path and tape device each take up a path.

Creating too many paths to your LUNs can potentially inhibit the total number of LUNs that can be allocated to an IBM i partition because there is also a limit of 64 LUNs per FC adapter *port* (direct or native connection) or virtual FC adapter (NPIV) port. NPIV can support up to 64 active and 64 passive LUNs per virtual path.

LUNs are only presented to the host from the owning IO group port.

You can map up to 64 virtual FC (NPIV) adapters from separate IBM i LPARs to the same FC port in VIOS.

This LUN limit only applies to IBM i clients as the limitation is enforced by the IBM i licensed internal code. The limitation of 64 partitions sharing a single FC port is enforced by the HMC/VIOS, and so that applies to any type of client partition.

Direct attached or NPIV mapped LUNs for IBM i require an “IBM i host kit” on the storage subsystem. On the Storewize V7000, IBM i uses the generic host support and does not require a specific host kit.

## **6.2 Performance**

When creating an open storage LUN configuration for IBM i as a client of VIOS, it is crucial to plan for both capacity and performance. As LUNs are virtualized for IBM i by VIOS instead of being directly connected it may seem that the virtualization layer will necessarily add a significant performance overhead. However, internal IBM performance tests clearly show that the VIOS layer adds a negligible amount of overhead to each I/O operation. Instead, the tests demonstrate that **when IBM i uses open storage LUNs virtualized by VIOS, performance is almost entirely determined by the physical and logical configuration of the storage subsystem.**

The IBM Rochester, MN, performance team has run a significant number of tests with IBM i as a client of VIOS using open storage. The resulting recommendations on configuring both the open storage and VIOS are available in the latest *Performance Capabilities Reference* manual (PCRM) at: <http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>. Chapter 6 focuses on virtualized storage for IBM i. In most cases, an existing IBM i partition using physical storage will be migrated to open storage LUNs virtualized by VIOS. The recommended approach here is to start with the partition's original physical disk configuration; then create a similar setup with the physical drives in the open storage subsystem on which LUNs are created, while following the suggestions in the PCRM sections. A basic rule of thumb is to make at least 6-8 LUNs for an IBM i partition.

The commonly used SAN disk sizing tool Disk Magic can also be used to model the projected IBM i performance of different physical and logical drive configurations on supported subsystems. You can work with IBM Techline or your IBM Business Partner for a Disk Magic analysis. The latest version of Disk Magic includes support for multiple open storage subsystems and IBM i as a virtual client of VIOS.

## **6.3 Dual hosting and multi-path I/O (MPIO)**

An IBM i client partition in this environment has a dependency on VIOS: if the VIOS partition fails, IBM i on the client will lose contact with the virtualized open storage LUNs. The LUNs would also become unavailable if VIOS is brought down for scheduled maintenance or a release upgrade. To remove this dependency, two or more VIOS partitions can be used to simultaneously provide virtual storage to one or more IBM i client partitions.

### **6.3.1 Dual VIOS LPARs with IBM i mirroring**

Prior to the availability of redundant VIOS LPARs with clients-side MPIO for IBM i, the only method to achieve VIOS redundancy was to use mirroring within IBM i. This configuration uses the same concepts as that for a single VIOS described in the “Virtual SCSI and Storage virtualization concepts” section. In addition, at least one additional VSCSI client adapter exists in the client LPAR, connected to a VSCSI server adapter in the second VIOS on the same Power server. A second set of LUNs of the same number and size is created on the same or a different

open storage subsystem, and connected to the second VIOS. The host-side configuration of the second VIOS mimics that of the first host, with the same number of LUNs (hdisks), vtscsiX and vhostX devices. As a result, the client partition recognizes a second set of virtual disks of the same number and size. To achieve redundancy, adapter-level mirroring in IBM i is used between the two sets of virtualized LUNs from the two hosts. Thus, if a VIOS partition fails or is taken down for maintenance, mirroring will be suspended, but the IBM i client will continue to operate. When the inactive VIOS is either recovered or restarted, mirroring can be resumed in IBM i.

### 6.3.2 Path redundancy to a single set of LUNs

Note that the dual-VIOS solution above provides a level of redundancy by attaching *two separate sets* of open storage LUNs to the same IBM i client through separate VIOS partitions. It is not an MPIO solution that provides redundant paths to a *single set* of LUNs. There are two MPIO scenarios possible with VIOS that remove the requirement for two sets of LUNs:

- A single VIOS partition using two FC adapters to connect to the same set of LUNs
- Two VIOS partitions providing redundant paths to the same set of LUNs on a single open storage subsystem

#### 6.3.2.1 Subsystem Device Driver – Path Control Module (SDDPCM)

VIOS also supports the Subsystem Device Driver – Path Control Module (SDDPCM) for certain storage subsystems. Examples of supported subsystems include the SAN Volume Controller (SVC), Storwize V7000 and IBM DS8000 system. To find out whether a particular storage system supports SDD-PCM for VIOS, refer to its interoperability matrix on the SDDPCM website: <https://www-304.ibm.com/support/docview.wss?uid=ssg1S4000201>. Note that there are separate support statements for AIX and VIOS. Always check if the VIOS version is supported by a specific SDDPCM driver version, and that you download the correct SDDPCM driver for the correct storage system and VIOS. If SDDPCM is supported on your storage subsystem for VIOS, download and install the driver following the instructions in the *Multi-path Subsystem Device Driver User's Guide* at the same location.

#### 6.3.2.2 Path redundancy with a single VIOS

If a VIOS LPAR has two or more FC adapters assigned and the correct host configuration is created in the open storage subsystem, VIOS will have redundant paths to the LUNs connected to the FC adapters. VIOS includes a basic MPIO driver, which has been the default instead of the RDAC (Redundant Disk Array Controller) driver since November 2008. The MPIO driver can be used with any storage subsystem which VIOS supports and is included in a default install. In this case, configuration is required only on the storage subsystem in order to connect a single set of LUNs to both ports on a FC adapter owned by VIOS.

This multi-path method can be configured in two ways: round-robin, or failover. For storage systems such as XIV, V7000 and DS8000 that support active-active multi-path connections, either method is allowed. Path failover will use one path and leave the other idle until the first path fails, while round-robin will always attempt to use both of the paths, if available, to maximize throughput. For systems that only support active-passive connections, such as IBM System Storage DS5000, DS4000, and DS3000 subsystems, failover is the only method allowed.

The access method can be configured on a per-LUN basis from the VIOS command line through a Telnet or SSH session. To show the current multi-path algorithm, run the following command:

- **lsdev -dev hdiskX -attr**
- Look for the **algorithm** attribute value. Also look at the **reserve\_policy** and **queue\_depth** attributes.



- When you have multiple paths to the same storage, the **reserve\_policy** should be set to **no\_reserve** to allow either path to access the storage.
- The **queue\_depth** tells VIOS how many concurrent I/O requests to send to the storage controllers. Check the documentation for the storage subsystem to determine the correct value to use. A large value can flood some of the entry level subsystems and cause problems.
- For SVC or Storwize V7000: The attribute **algorithm** should be set as follows:
  - If the driver **SDDPCM** is used in VIOS, the attribute should be set to **load\_balance**,
  - If the **MPIO** (built in) driver is used the attribute should be set to **round\_robin**.

To change the multi-path algorithm, run the following command:

- **chdev -dev hdiskX -attr algorithm=round\_robin**, or
- **chdev -dev hdiskX -attr algorithm=fail\_over**

These commands must be repeated for each hdisk.

### 6.3.2.3 Redundant VIOS LPARs with client-side MPIO (VSCSI)

Beginning with IBM i 6.1 with Licensed Internal Code (LIC) 6.1.1, the IBM i VSCSI client driver supports MPIO through two or more VIOS partitions to a single set of LUNs (up to a maximum of eight VIOS partitions). This multipath configuration allows a VIOS partition to fail or be brought down for service without IBM i losing access to the disk volumes as the other VIOS partition(s) remain active. Figure 9 illustrates the new capability for IBM i.

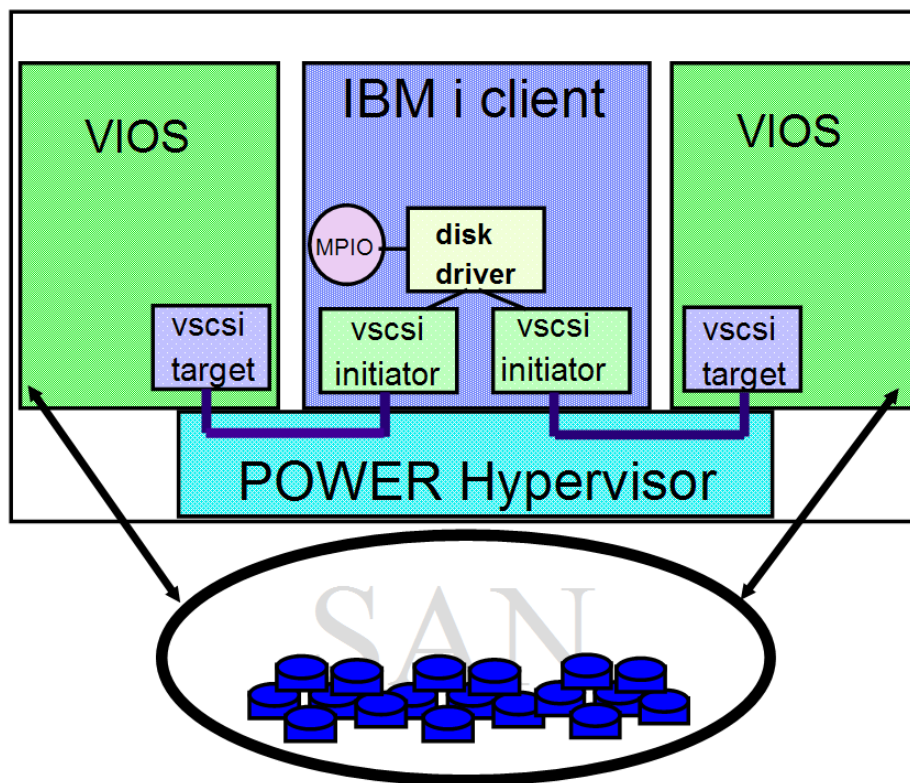


Figure 9: Redundant VIOS server concept.

Note that as with dual VIOS LPARs and IBM i mirroring, a VSCSI connection from IBM i to *both* VIOS partitions is required. On the storage subsystem, it is recommended to create a host for each VIOS containing the worldwide port names (WWPNs) of all FC adapters in that LPAR, and then create a host group comprising those VIOS hosts. A single set of LUNs for IBM i can then be created and mapped to the host group (shared access), giving both of the VIOS LPARs access to the LUNs. Configuring redundant VIOS LPARs with client-side MPIO does not preclude using MPIO within a VIOS for a higher level of resiliency, as described in the previous section.

For storage systems such as DS8000 and XIV, these connections are typically shared through the round-robin access as described earlier. So it is a good practice to have two or more connections to each VIOS partition.

Note: For D5000, DS4000 and DS3000 storage subsystems with dual controllers, a connection must be made to both of the controllers to allow an active and a failover path. When the volumes are created on these systems, the host OS type should be set to *DEFAULT* or *AIX* (not *AIX ADT/AVT*, or failover/failback oscillations might occur). For all storage systems, it is recommended that the fabric configuration uses **separate dedicated zones** and FC cables for each connection.

After the VSCSI client driver in IBM i detects a second path to the same set of LUNs through a different VIOS LPAR, the disk names change from DDxxx to DMPxxx. This is identical to multi-path I/O to a directly attached DS8000 subsystem. If a path is lost, the disk names do not change back to DDxxx, but a path failure message is sent to the QSYSOPR message queue in IBM i. As soon as the path is restored, an informational message is sent to the same queue. To monitor the status of the paths in IBM i:

- Start Services Tools and sign in.
- Work with disk units.
- Display disk configuration.
- Display disk path status.

Note that while the HMC allows a single VSCSI client adapter to be tagged as an IPL device, it is not required to change the tagged resource if that path is lost. As long as the alternate path through the second VIOS is active, the IBM i LPAR will be able to IPL from the same load source LUN. The alternate path must exist prior to the loss of the original IPL path and prior to powering the IBM i LPAR off.

Another option to VSCSI storage, as just described, is NPIV attached storage facilitated through both VIOS partitions. The sets of commands listed earlier in this section do not come into play because the VIOS partitions become pass thru virtual servers.

## **6.4 Best practices for Fibre Channel storage**

When configuring LUNs for IBM i (virtualized by VIOS), follow the best practices outlined in chapter 18 of the latest *Performance Capabilities Reference* manual, available here: <http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>. Note that some of its recommendations apply only to IBM i using virtual storage outside of the chassis environment.

In addition to the guidelines in the *Performance Capabilities Reference* manual, follow these additional recommendations:

- Use Fibre Channel disk drives (and not SATA or FATA) to create the RAID ranks/arrays for production IBM i workloads
- Use 15K RPM drives for medium and heavy I/O IBM i workloads, and 10K RPM drives for low I/O workloads

- When creating a host connection to the WWPN of the Fibre Channel card on the node, specify at most two specific host ports. Do not create the connection so that the Fibre Channel adapter can connect to all host ports on the storage subsystem, which is the default for some subsystems.
- Properly zone the switches between the adapter ports and the SAN host ports.
- When configuring vSCSI LUNs for IBM i within VIOS: as soon as the LUNs that will be virtualized to IBM i have reported in VIOS, and prior to assigning them to the IBM i client, change their queue depth to improve performance. Start a Telnet session to VIOS and login as **padmin**. Use the following command for each LUN = hdisk:
  - **lsdev -attr -dev hdisk#**, where # is the hdisk number you want to display.
  - **chdev -perm -attr queue\_depth=N -dev hdisk#**
    - **Note:** The order of the parameters can be changed, as shown, to facilitate repeating the command and only having to alter the hdisk number.
    - **Note:** The recommended value for N on the queue\_depth parameter will vary depending on the type of external storage being used. For example, a good value to start with for DS8000, V7000, and XIV is 32. Low-end storage servers such as DS3000 may perform better when queue\_depth is set to 8. If I/Os are being queued on the VIOS, you may need to reduce this value to optimize performance.
    - **Note:** Some of the low end SANs might not handle the larger number of concurrent commands as well, which can adversely affect performance.
    - For redundant VIOS servers, each server needs to be able to access the same hdisks, so another attribute needs to be set for this: **reserve\_policy=no\_reserve** on each hdisk. Add a space between the attribute lists on the command.
  - **lsdev -attr -dev hdisk#**, to validate the change.
- Another parameter that can improve performance is the number of I/O commands to send to the Fibre Channel adapter. The recommended value is 1024. Be sure to change the value for all ports of all of the FC adapters:
  - **lsdev -attr -dev fcs#**, where # is the FC adapter number you want to display.
  - **chdev -attr num\_cmd\_elems=1024 -perm -dev fcs#**
  - **lsdev -attr -dev fcs#**, where # is the FC adapter number you want to display to validate the change.
- To improve reliability, enable dynamic tracking and fast fail over for the LUNs virtualized to IBM i. Fast\_fail should be done only when multiple paths to the disks exist. Do so for all ports of the all of the adapters:
  - **lsdev -attr -dev fscsi#**, where # is the FC adapter number you want to display.
  - **chdev -attr dyntrk=yes fc\_err\_recov=fast\_fail -perm -dev fscsi#**
  - **lsdev -attr -dev fscsi#**, where # is the FC adapter number you want to validate.

Using the **-perm** option in these commands means that the value will be updated only in the VIOS device configuration database (ODM). To make the change effective, **reboot VIOS** when there is downtime available for all client partitions.

These settings are applied to each of the fscsiX devices being used. To show the current fscsiX devices, run the following command:

- **lspath**

Set the fast fail and dynamic tracking attributes on these devices, run the following command:

- **chdev-attr fc\_err\_recov=fast\_fail,dyntrk=yes -perm -dev fscsiX**

Next, restart VIOS. The chdev command must be repeated for each fscsiX device but only one restart is required after all fscsiX attributes have been set.

NOTE: fast\_fail is designed for a multipath configuration to not retry failed paths for long periods of time. In a single path configuration, do not set fast\_fail since it will not tolerate temporary failures on the single path.

To verify the settings have been changed successfully, run the following command:

- **lsdev -dev fscsiX -attr fc\_err\_recov**
- **lsdev -dev fscsiX -attr dyntrk**
- Or just **lsdev -dev fscsiX -attr** and review the list of attributes shown.

Importantly, the SCSI reserve policy for *each LUN (or hdisk) on both VIOS LPARs* must be set to *no\_reserve* to allow disk sharing. Some storage subsystems such as XIV default to *no\_reserve* and do not require a change, while others such as DS4000 and DS5000 default to *single\_path*.

NOTE: The change must be made **prior** to mapping the LUNs to IBM i and it does not require a restart of VIOS.

To show the current reserve policy settings, run the following command:

- **lsdev -dev hdiskX -attr reserve\_policy**

To **set** the *no\_reserve* attribute if necessary:

- **chdev -dev hdiskX -attr reserve\_policy=no\_reserve**

This is done for every hdiskX

## 7 Attaching open storage to IBM i through VIOS

As described in the “IBM i hosted by VIOS

Virtual SCSI and Storage virtualization concepts” section, IBM i joins the VIOS virtualization environment, allowing it to use open storage. The setup process involves the following steps:

- Open storage configuration
- VIOS installation and configuration
- SAN zoning
- IBM i installation and configuration

The VIOS and open storage configuration steps are the same as for existing clients of VIOS, such as AIX and Linux. These steps are well documented and any existing skills in those areas would apply to IBM i as a client partition, as well. The only significant difference with IBM i is the specific open storage configuration requirements in order to achieve good performance, as referenced in the “” section.

### 7.1 Open storage configuration

All open storage subsystems follow similar general configuration steps:

- Perform physical disk and Redundant Array of Independent Disks (RAID) configuration
- Create volumes (LUNs)
- Attach those volumes to a host (VIOS in this case)

Naturally, the detailed steps and graphical interfaces used vary by subsystem. Locate the correct documentation for your supported subsystem in IBM Redbooks at: <http://www.redbooks.ibm.com>. Importantly, follow the instructions for creating and attaching LUNs to an AIX host if you are planning to use VSCSI adapters. For virtual FC the IBM i host type is used. Refer to the associated sections for more information.

### 7.2 VIOS installation and configuration

Refer to the VIOS planning, installation and configuration topics in the *PowerVM Editions Guide*, available at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm>.

As described in the “” section, as soon as VIOS is installed and physical I/O resources are assigned to it, those resources can be virtualized to IBM i using the HMC/SDMC. It is typically not necessary to use the IVM/VIOS command line to assign virtual disk, optical and network resources to IBM i.

As you add client workload behind VIOS, consider increasing the VIOS CPU and memory allocations. Using a shared processor pool with uncapped resources for the VIOS CPU is recommended. When using high speed adapters (10Gb Ethernet or 8Gb FC) the amount of memory used by the VIOS partition will increase. Increase the memory allocation to the VIOS partitions to 6-8GB.

### 7.3 SAN Zoning

Fibre channel SAN switches can be zoned using their CLI or GUI. The GUI's are similar across the vendors, but the terms used are different. You create an Alias name that is associated with a WWPN from a host adapter on a server or a host port on the SAN controller. The aliases are then combined into sets comprised of a host WWPN and the SAN host port's WWPNs that the host can access. These zone sets are combined under a zone configuration for saving and enabling/activating. See the switch manuals for the specifics.

To assist in the SAN zoning of the virtual WWPNS associated with NPIV configurations there are management console commands to list and login/logout the virtual WWPNS. The login option will make them appear on the FC switches to enable easier configuration. See the `lsnportlogin`, `chnportlogin` commands in the HMC command reference located here:

<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha1/usingthehmcremotecommandline.htm>

Starting with V7.7.5 of the HMC, there is a browser interface to log into the virtual fibre channel adapters.

- On the HMC navigation pane, select the system and your IBM i LPAR
- In the lower menu, expand **Configuration** and then **Managed Profiles**
- Select the profile and click **Edit** from the Action pulldown menu.
- Click on the **Virtual Adapters** tab and select the Client Fibre Channel adapter.
- From the Actions pulldown, select **Advanced -> Login Logout Fibre Channel**
- On the following window, click the **Login** button

The WWPNS status should change to 1 and VIOS should now be logged into the adapter. At this point, you can configure the switch and storage server using the virtual WWPNS that appear on the Login window.

## 7.4 Mapping storage to IBM i using the HMC

Use the following steps:

- On the HMC navigation pane, click on **Managed Systems**. Then select the Power server by clicking in the **checkbox** next to it.
- In the lower menu, expand **Configuration** and then **Virtual Resources**.
- Click **Virtual Storage Management**.
- Select the correct VIOS from the list and click **Query VIOS**.
- Click the **Physical Storage** tab.
- A list of **hdisks** is shown.
- In the upper middle section of the pane, use the pull down to choose the combination of **partition name** and **virtual SCSI adapter** to associate the hdisk to. You can map up to 16 hdisks to each VSCSI adapter (i7.1 TR8 this limit has been increased to 32), but the interface will not stop you from mapping more. This is because AIX partitions can also use this interface and AIX supports up to 256 hdisks per adapter (though they seldom map that many). If you map more than 16/32, the additional hdisks will not be seen by IBM i.
- Choose an unassigned hdisk to map to the adapter and click **Assign**.
- Repeat this process for each hdisk. As the mapping completes, the IBM i client partition should be listed as the new hdisk owner.
- **Close** the window when done.

## 7.5 Mapping storage to IBM i using the SDMC

Use the following steps:

- On the SDMC welcome page, click on **Hosts**.
- Right click on the Power blade and select **System Configuration -> Virtual Resources -> Virtual Storage Management**.
- Select the correct VIOS from the list and click **Query VIOS**.
- Click the **Physical Storage** tab.
- A list of **hdisks** is shown.
- In the upper middle section of the pane, use the pull down to choose the combination of **partition name** and **virtual SCSI adapter** to associate the hdisk to. You can map up to 16 hdisks to each VSCSI adapter (i7.1 TR8 this limit has been increased to 32), but the interface will not stop you from mapping more. This is because AIX partitions can also

use this interface and AIX supports up to 256 hdisks per adapter (though they seldom map that many). If you map more than 16/32, the additional hdisks will not be seen by IBM i.

- Choose an unassigned hdisk to map to the adapter and click **Assign**.
- Repeat this process for each hdisk. As the mapping completes, the IBM i client partition should be listed as the new hdisk owner.
- **Close** the window when done.

## 7.6 NPIV Configuration for IBM i LUNs or tape media libraries

Note: If you add virtual FC adapters using the Dynamic LPAR operation, use the save configuration to the partition profile option when done. If you instead edit the profile separately, you will generate a new pair of virtual WWPNS.

There are three general steps in configuring NPIV for IBM i:

- LPAR and VIOS configuration on the Power server's management console
- Storage subsystem or tape library configuration
- SAN zoning

To perform the LPAR and VIOS setup, refer to Chapter 2.9 in the Redbooks *PowerVM Virtualizing Managing and Monitoring* (SG 24-7590) at:

<http://www.redbooks.ibm.com/abstracts/sg247590.html?Open>. While the examples given are for an AIX client of VIOS, the procedure is identical for an IBM i client.

To perform the storage or tape configuration, refer to the Redbooks *IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i* (SG24-7120) at:

<http://www.redbooks.ibm.com/abstracts/sg247120.html?Open> or *Implementing IBM Tape in i5/OS* (SG 24-7440) at <http://www.redbooks.ibm.com/abstracts/sg247440.html?Open>. As mentioned above, from the storage subsystem's or tape library's perspective, the configuration is identical to that for IBM i directly attached through the SAN fabric.

There is a web interface on the tape media library where you need to enable control paths from each device that you want IBM i to be able to work with. Selecting Enable, creates the control paths. IBM i cannot dynamically detect these control paths. To detect the control paths, you need to re-IPL the virtual I/O adapter (IOA).

First determine which virtual IOA has been created for the virtual FC adapters. To do this, enter a WRKHDWRSC \*STG command and check for a 6B25 (virtual FC) adapter. Note the IOP/IOA name.

Next, use STRSST command:

- Use option 1 to **Start a service function**
- Option 7 **Hardware service manager**
- Option 2 **Logical Hardware Resources**
- Option 1 **System bus resources**
- Enter a **255** in the **System bus(es)** to work with field and hit **enter**. This is the virtual adapter bus.
- Locate the virtual IOA from above
- Enter an option 6 for **I/O debug**
- Then option 4 to **IPL the IOA**.
- Use **F3** to exit SST.

Return to the WRKHDWRSC \*STG and use an option 9 to Refer to the tape devices under the VFC IOA. With auto configuration turned on, a new tape device(s) should show up under WRKCFGSTS \*DEV TAP\*.

You may have to check the **port type** defined on the tape media library for the fibre channel port associated with the tape device.

- Log into the tape library interface
- Go to **Configure Library**
- Then select **Drives**
- Set the port type to **N-Port**

Accordingly, DS8000 LUNs might be created as “IBM i protected” or “IBM i unprotected” and will correctly report as such to Storage Management in IBM i. A tape library and a drive within will also report the correct device names and types, such as TAPMLBxx, 3584 and TAPxx, 3580-004 and so on. All tape operations supported with direct FC attachment are supported through NPIV, including hardware tape encryption.

## 8 IBM i tape options

An IBM i client partition with a VIOS host can use a mix of virtual and physical I/O resources.

If the VIOS hosting partition supports a single physical tape device, it can be virtualized to the client partition(s) by using a VSCSI adapter pair between the VIOS and the client IBM i. This is facilitated by the management console being used. See the following sections for details.

For the current list of supported tape devices and associated operating system levels please refer to the DevelopWorks web page for IBM Removable Media on IBM i:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/IBM%20Removable%20Media%20on%20IBM%20i>

VIOS also supports virtualization of a tape media library through the NPIV configuration described in section “NPIV Configuration for IBM i LUNs or tape media libraries”.

For the current list of supported tape libraries and associated operating system levels please refer to the DevelopWorks web page for IBM Removable Media on IBM i:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/IBM%20Removable%20Media%20on%20IBM%20i>

Another approach is to assign a physical tape adapter directly to the IBM i partition. The tape adapter can be any adapter supported by IBM i on IBM Power servers and can be configured to be shared with other partitions.

To assign a physical tape adapter to the IBM i client partition, refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.

### 8.1 Virtual tape configuration using the HMC

There is not a graphical interface for the configuration of virtual tape at this time. Perform the following steps to configure virtual tape using the VIOS CLI:

- Use telnet or PuTTY to connect to the VIOS partition. Sign in using **padmin** as the user ID.
- Enter **cfgdev** to check for new devices.
- Enter **lsdev | grep rmt** to view the tape devices and ensure that they are in **Available** state.



- Enter **lsdev | grep vhost** and note the last vhost listed there.
- You need to associate this device with a VSCSI adapter pair. You need to use the HMC interface to create those. Refer to the “Creating multiple Virtual SCSI adapters using the HMC” section for details. Then return to this step.
- On the VIOS CLI enter **lsdev | grep vhost**. There should be a new vhostY listed. This vhostY is the VSCSI adapter in VIOS that you just created.
- To map the tape drive to the vhostY, enter **mkvdev -vdev <rmtX> -vadapter vhostY**.
- Enter **lsmmap -all | more** and press the **Spacebar** key to advance through the mappings. Look for the **vhostY** and make sure the **vttapeZ** device is associated with it.
- On the IBM i virtual server with auto configuration turned on, a TAPxx device appears. Vary it on to use it.

## 8.2 Virtual tape configuration using the SDMC

There is not a graphical interface for the configuration of virtual tape at this time. Perform the following steps to configure virtual tape using the VIOS CLI:

- Use telnet or PuTTY to connect to the VIOS partition. Sign in using **padmin** as the user ID.
- Enter **cfgdev** to check for new devices.
- Enter **lsdev | grep rmt** to view the tape devices and ensure that they are in **Available** state.
- Enter **lsdev | grep vhost** and note the last vhost listed there.
- You need to associate this device with a VSCSI adapter pair. You need to use the SDMC interface to create those. Refer to the “Creating multiple Virtual SCSI adapters using the SDMC” section for details. Then return to this step.
- On the VIOS CLI enter **lsdev | grep vhost**. There should be a new vhostY listed. This vhostY is the VSCSI adapter in VIOS that you just created.
- To map the tape drive to the vhostY, enter **mkvdev -vdev <rmtX> -vadapter vhostY**.
- Enter **lsmmap -all | more** and press the **Spacebar** key to advance through the mappings. Look for the **vhostY** and make sure the **vttapeZ** device is associated with it.
- On SDMC: update the inventory to view the changes.
- On the IBM i virtual server with auto configuration turned on, a **TAPxx** device appears. Vary it on to use it.

## 8.3 Virtual tape configuration using IVM

VIOS will see a supported tape device as an **rmt#** device. If you physically connected the tape device to its associated adapter *after* VIOS came up, you will need to use the **Hardware** link and click on the **Configure Devices** button to see the rmt# device.

IVM has an interface to move the tape device under the **View/Modify Storage** pane

- Click the **Tape** tab.
- Select the **rmt#** device and assign it to the IBM i partition.
- If auto configuration is enabled on the IBM i partition, the tapxx device will be configured.
- Vary on the device using the **WRKCFGSTS \*DEV TAP\*** interface.

## 9 IBM i installation and configuration

The IBM i client partition configuration as a client of VIOS is the same as that for a client of an IBM i 6.1 host partition. Refer to the **Creating an IBM i logical partition that uses IBM i virtual I/O resources using the HMC** topic in the *Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>. The guide does not address VIOS hosting IBM i, but the VSCSI and virtual Ethernet adapter concepts are the same.

For SDMC use the IBM Systems Director Management Console: Introduction and Overview in IBM Redbooks at: <http://www.redbooks.ibm.com/abstracts/sq247860.html?Open>

### 9.1 Post-IBM i install tasks and considerations

#### 9.1.1 Configure IBM i networking

After the IBM i client partition is installed and running, the first system management step is to configure networking. There are three types of network adapters that can be assigned to an IBM i client partition:

- A standard physical network adapter in a PCI slot
- A logical port on a host Ethernet adapter (HEA)
- A virtual Ethernet adapter

Note that both physical and virtual I/O resources can be assigned to an IBM i virtual client partition. If a physical network adapter was not assigned to the IBM i client partition when it was first created, Refer to the topic **Managing physical I/O devices and slots dynamically using the HMC** in the *Logical Partitioning Guide* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>) to assign an available adapter.

An IBM i client partition can also use the HEA capability of certain POWER6 processor-based servers. To assign a logical port (LHEA) on an HEA to an IBM i client partition, Refer to the **Creating a Logical Host Ethernet Adapter for a running logical partition using the HMC** topic in the *Logical Partitioning Guide* at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.

A virtual Ethernet adapter can also provide network connectivity to an IBM i client partition. To create one, refer to the topic **Configuring a virtual Ethernet adapter using the HMC** in the *Logical Partitioning Guide*: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.

For the SDMC interfaces for networking refer to the IBM Systems Director Management Console: Introduction and Overview Redbook at: <http://www.redbooks.ibm.com/abstracts/sq247860.html?Open>

In all cases, the assigned network adapter will be recognized as a communications port (CMNxx) in IBM i. The type of communications port will depend on the network adapter used: for example, 5706 for a Gigabit Ethernet adapter, 5623 for an LHEA and 268C for a virtual Ethernet adapter. In the case of a standard PCI network adapter or an LHEA, networking can be configured following the process described in the IBM i networking topic in the Information Center at: <http://publib.boulder.ibm.com/infocenter/iseriess/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm>.

A virtual Ethernet adapter is recognized by IBM i as a communications port (CMNxx) of type 268C:

```

Work with Communication Resources
System: Y213B2P2
Type options, press Enter.
5=Work with configuration descriptions 7=Display resource detail

Opt Resource Type Status Text
- CMB03 268C Operational Comm Processor
- LIN02 268C Operational LAN Adapter
- CMN02 268C Operational Ethernet Port

```

Figure 10 Sample communications resources.

If the IBM i client partition is to use a virtual Ethernet adapter to communicate with an external network, additional configuration in the VIOS hosting partition must be done. An SEA must be created in VIOS to bridge the internal virtual LAN (VLAN) to the external LAN. Use the HMC/SDMC and the instructions in section Network virtualization to perform the SEA configuration. Ensure that the port VLAN ID (PVID) that you choose for the client IBM i's VE adapter matches the PVID that VIOS has in its VE adapter.

### 9.1.2 End-to-end LUN device mapping

On October 2009, IBM enhanced both the HMC and VIOS to allow end-to-end device mapping for LUNs assigned to client LPARs, such as IBM i. The new function enables administrators to quickly identify which LUN reporting in VIOS (or, hdisk) is which DDxxx disk device in IBM i. This in turn makes it easier to troubleshoot disk-related problems and safer to change a virtualized disk configuration. In order to correctly perform the mapping, the HMC requires an active RMC connection to VIOS.

To perform end-to-end LUN device mapping, use the following steps:

- Sign in to the HMC as **hscroot** user ID.
- Expand **Systems Management**.
- Expand **Servers**.
- Click the correct managed system (server).
- Select the correct VIOS by using the checkbox.
- Click **Hardware Information** → **Virtual I/O Adapters** → **SCSI**.
- You will be shown a list of drives associated with their VIOS hdisks. Note the DD## and its associated hdisk# that you are interested in.
- Click back on the word **Servers** in the left hand navigation pane of the HMC.
- Select the correct managed server by selecting its checkbox.
- In the menu below, expand **Configuration** and then **Virtual Resources**.
- Click **Virtual Storage Management**.
- Select the correct VIOS from the list and click **Query VIOS**.
- Click the **Physical Volumes**.
- The hdisks that VIOS sees are shown along with the partition that they are assigned to. On the far right side of the **Physical Location Code** column, there is a -L#00000... This is the LUN number associated with the hdisk. This is a hexadecimal number.
- Use the SAN interface to determine which volume has that LUN number. You may have to convert the hex number to a decimal number (I know, it's been a while, but you can do it!). If the SAN is the V7000, look for the SCSI ID as the LUN #.

### 9.1.3 How to perform IBM i operator panel functions

If the system is HMC-managed, follow these steps:

- Sign in to the HMC with a profile with sufficient authority to manage the IBM i client partition
- Select the partition.
- Use the open-in-context arrow to select **Serviceability** → **Control Panel Functions** and then the desired function.

If the system is IVM-managed, follow these steps:

- In IVM, click **View/Modify Partitions** .
- Select the IBM i partition.
- From the **More Tasks** list, select **Operator panel service functions**
- Select the function you wish to perform and click **OK**

If the system is SDMC-managed, perform follow these steps:

- On the SDMC welcome page, select the host you are working with.
- Right-click the host and click **Service and Support-> Control Panel Functions-> (20) Type, Model, Feature**. The function is limited to displaying the hardware information listed.

### 9.1.4 Displaying the IBM i partition System Reference Code history

If the system is HMC-managed, follow these steps:

- Sign in to the HMC with a profile that has sufficient authority to manage the IBM i client partition.
- Select the partition.
- Use the open-in-context arrow and then select **Serviceability → Reference Code History**.
- To display words 2 through 9 of a reference code, select the option for that code.

If the system is IVM-managed, perform the following steps:

- In IVM, click **View/Modify Partitions**.
- Select the IBM i partition.
- From the **More Tasks** drop-down menu and select **Reference Codes**.
- Click an SRC to display all words.

If the system is SDMC-managed, perform the following steps:

- View the blade's system reference codes from the SDMC welcome page.
- Select the host for the Power server
- Right-click the host and click **Service and Support-> Reference Code history**.

### 9.1.5 Client IBM i LPARs considerations and limitations

Refer to the **Considerations and limitations for i5/OS client partitions on systems managed by the Integrated Virtualization Manager (IVM)** topic in the Information Center at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzahc/rzahcbladei5limits.htm>.

### 9.1.6 Configuring Electronic Customer Support (ECS) over LAN

A supported WAN adapter can be assigned to the IBM i client partition for ECS. Alternatively, ECS over LAN can be configured. Refer to the **Setting up a connection to IBM** topic in the Information Center at:

[http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzaji/rzaji\\_setup.htm](http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm).

### 9.1.7 Backups

IBM i as a client of VIOS can use a mix of virtual and physical I/O resources. One approach is to assign an available tape adapter on the system to the IBM i client partition and treat it as a standard IBM i partition. The tape adapter can be any adapter supported by IBM i on IBM Power servers and can be shared with other partitions. To assign an available tape adapter to the IBM i client partition, refer to the **Managing physical I/O devices and slots dynamically using the HMC** topic in the *Logical Partitioning Guide* at:

<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>.

After a tape adapter connected to a tape drive or library is available to the client partition, refer to the Backup and Recovery topic in the Information Center to manage backups:

[http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int\\_215989](http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahg/rzahgbackup.htm&tocNode=int_215989)

## 10 DS5000 direct attached storage to IBM i

### 10.1 Overview

IBM i has been able to connect to 512-byte-per-sector open storage through VIOS since early 2008 and the solution is implemented in many production environments. On October 2009, IBM further simplified open storage use for IBM i by announcing direct FC attachment to the DS5100 and DS5300 subsystems. With the new enhancement, the IBM i LPAR owns one or more FC adapters, which are connected to the SAN fabric. 512-byte LUNs on the DS5100 or DS5300 are then mapped to the world-wide port names (WWPNs) of the FC adapter(s) in the IBM i LPAR. VIOS is no longer required in this case. Note that one or more VIOS LPARs may still exist on the same Power server to take advantages of other technologies, such as Active Memory Sharing (AMS), NPIV, Live Partition Mobility or other I/O virtualization. However, VIOS is no longer required in order for IBM i to use storage from DS5100 and DS5300 system. For example, it is possible to have the same or another IBM i LPAR on the same Power server and still use VIOS for optical virtualization.

A note on terminology: while the new enhancement is referred to as **native attachment of IBM i to DS5100 and DS5300**, IBM expects that in most cases the attachment will be through a FC switch. While directly attaching FC cables from the adapters on the Power server owned by IBM i to the host channel ports on the storage is supported, this approach decreases the flexibility of the SAN environment. Therefore, **native attachment** in this case refers to **without VIOS**.

There are two main benefits of natively attaching IBM i to DS5100 and DS5300 systems:

- VIOS skills are not required to implement and use the solution, only IBM i and SAN skills are required.
- The processor, memory, I/O and software expenditure to run a VIOS LPAR is not required

Multipath I/O for IBM i in this environment is supported with a maximum of eight paths. No additional software on the DS5100 or DS5300 systems or IBM i is required. The IBM i FC adapter driver employs the round-robin algorithm by default; it cannot be changed.

### 10.2 Supported hardware and software

As of the publication date, only the following hardware is supported for IBM i directly attaching to DS5100 and DS5300 systems. This support is also reflected in the “IBM i using open storage configurations” section.

- DS5100 (machine-type 1818, model 51A)
- DS5300 (machine-type 1818, model 53A)
- EXP810 expansion unit
- EXP5000 expansion unit
- EXP5060 expansion unit
- FC and SATA HDDs (hard disk drives), with FC drives recommended for production workloads
- POWER6 processor-based or higher Power servers
- FC #5273 4Gb dual-port PCI-e FC adapter
- FC #5276 8Gb dual-port PCI-e FC adapter
- FC #5774 4Gb dual-port PCI-e FC adapter
- FC #5735 8Gb dual-port PCI-e FC adapter

As of the publication date, only the following software or firmware is supported for IBM i directly attached to DS5100 and DS5300 systems. This support is also reflected in “IBM i using open storage configurations” section.

- IBM i 6.1 with LIC 6.1.1 or higher
- Controller firmware 07.60.28.00 or higher
- DS Storage Manager 10.60.x5.17 or higher
- IBM i Host Attachment Kit, FC #7735 (required)
- Storage Partitioning (strongly recommended)

### **10.3 Best practices, limitations and performance**

As with other storage subsystems, it is recommended to create LUNs for IBM i production workloads on FC drive RAID arrays only in order to meet the performance requirements of the application. While SATA drives are supported for direct attachment to IBM i, they are best suited for test, development and archival of IBM i workloads. Mixing FC and SATA drives within an enclosure is supported along with solid-state drives (SSDs). Hardware encryption of the data on the physical drives within the storage subsystem is supported.

All LUNs mapped directly to IBM i must be protected. Therefore, the supported RAID levels for IBM i direct attachment are 1, 5, 6 and 10. LUNs report as **protected** in IBM i; therefore, IBM i mirroring is not possible. The reliability of the solution is based on the RAID protection within the storage subsystem and on using multiple paths. The maximum LUN size for IBM i in this environment is up to but **not** including 2TB (remember the load source minimum and maximum sizes). The maximum number of LUNs per FC adapter port is 64. See the “Storage planning” section for a discussion on *paths* and their limits. Dynamic volume (LUN) expansion is not supported for IBM i. After LUNs have been mapped to IBM i, at most 300 from the same subsystem can be added to an ASP with a single operation. More than 300 LUNs can be added to an ASP in general, but they require multiple add operations.

The requirement for the load source LUN size is based on the operating system level being installed. In i7.1 and beyond the load source must be at least 70 GB, but to be safe make it 80 GB. The load source does not need to be the very first LUN mapped to IBM i. When performing an Initial Program Load (IPL), an active path to the load source LUN is required. Path failover is not possible during IPL; however, as soon as the IPL is complete, failover functionality resumes.

As with other host types, each port on a FC adapter in IBM i should be zoned separately to controller A or controller B on the storage subsystem. Furthermore, it is strongly recommended to use Storage Partitioning when directly connecting to IBM i. Storage Partitioning is a premium feature on the DS5100 and DS5300 which allows a set of LUNs to be associated only with specified hosts and host ports. Storage Partitioning allows the creation of multiple host groups; only the hosts within a particular host group are allowed access to a group of LUNs. Without that premium feature, all hosts reside in the default host group, which allows IBM i to access LUNs mapped to other hosts. The recommended approach is to use Storage Partitioning and create a host group for each IBM i LPAR using the storage subsystem, placing the hosts associated with that LPAR in the host group. Hosts representing different IBM i LPARs should not be placed in the same host group, as IBM i does not support shared simultaneous access to the same LUNs from multiple LPARs.

Note that migrating from IBM i using DS5100 or DS5300 storage system through VIOS to IBM i using the same subsystem directly requires data migration, typically by save and restore. There is no option to reassign the LUNs from a VIOS host to an IBM i host or to use IBM Tivoli® Storage FlashCopy® Manager or VolumeCopy due to sector size differences..

Performance for IBM i directly attached to DS5100 and DS5300 system is very similar to that of the same adapter and storage configuration attached through VIOS. Direct attachment

performance will not be significantly better than that of VIOS attachment for IBM i. Performance results for IBM i directly attached to DS5300 system are in Chapter 5.1.2 of the *Performance Capabilities Reference* manual at: [http://www-03.ibm.com/systems/resources/systems\\_power\\_software\\_i\\_perfmgmt\\_pcrm\\_apr2011.pdf](http://www-03.ibm.com/systems/resources/systems_power_software_i_perfmgmt_pcrm_apr2011.pdf).

## 10.4 Sizing and configuration

There are two main sources of sizing information when planning a configuration involving IBM i directly accessing DS5100 or DS5300 systems: the Disk Magic sizing tool and the *Performance Capabilities Reference* manual mentioned in the previous chapter. It is recommended to use them jointly. Start with Disk Magic to determine the physical storage configuration, such as number of physical drives, number and type of RAID arrays and amount of cache. At the same time, the *Performance Capabilities Reference* manual offers some excellent suggestions on the number of LUNs and paths to use based on the I/O operations per second and throughput required. Note that the maximum number of LUNs for IBM i direct attachment is 64.

Configuration of IBM i direct attachment to DS5100 and DS5300 systems can be divided into two general steps:

- Storage and SAN zoning configuration
- Power server hardware and LPAR configuration

To configure the DS5100 or DS5300 systems, refer to the *IBM Midrange System Storage Hardware Guide* (SG24-7676) from IBM Redbooks available at: <http://www.redbooks.ibm.com/abstracts/sg247676.html?Open>, while following the best practices and limitations outlined in the previous section of this paper. When creating a host for direct IBM i attachment on the storage subsystem, the host type should be IBM i. To configure the SAN fabric, follow the instructions from your FC switch manufacturer.

To create the IBM i LPAR and assign the FC adapters to it, follow the instructions in the *Logical Partitioning Guide*, available at: <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphat/iphat.pdf>.

For SDMC refer to the IBM Redbooks IBM Systems Director Management Console: Introduction and Overview at: <http://www.redbooks.ibm.com/abstracts/sg247860.html?Open>

## 11 Redundant VIOS Virtual Servers hosting IBM i

For a production IBM i virtual server hosted by VIOS, you want redundant VIOS virtual servers in case of failures or scheduled maintenance on the hosting VIOS virtual server.

IVM cannot be used since it is actually running in VIOS and can only manage that virtual server. For a Power server or a POWER blade, either an HMC or an SDMC can be used. You need to consider the following steps for the configuration of this environment.

- Create two VIOS virtual servers/LPARs using SDMC/HMC. Allocate processor, memory, network adapters and storage adapters for each. For SDMC use the IBM Systems Director Management Console: Introduction and Overview from IBM Redbooks at: <http://www.redbooks.ibm.com/abstracts/sg247860.html?Open>
- Create an IBM i virtual server/LPAR using SDMC/HMC. Allocate processor and memory. The networking and storage are typically virtualized, but physical adapters can be assigned. If you are considering moving to IBM i hibernation and eventually to LPM, you should make the virtual server/LPAR fully virtualized hardware and host it with VIOS.
- You need shared storage for IBM i that both VIOS virtual servers/LPARs can access. This implies a SAN as integrated disks cannot share access. Select either a VSCSI or virtual FC configuration for the shared storage.
- For networking considerations, see section “Redundant VIOS Networking”.

## **12 Copy Services and IBM i**

### ***12.1 General Support Information***

Much support has been added in IBM i 6.1 and 7.1 for external storage devices.

IBM PowerHA SystemMirror for i provides an automated, IBM i-driven replication solution that allows clients to leverage their existing IBM i skills. PowerHA SystemMirror for i supports



FlashCopy, Metro Mirror, and Global Mirror with native and/or direct attach storage, with storage attached through VIOS, and with storage attached via NPIV. IBM PowerHA SystemMirror for i is supported by the IBM i Support Center. For more information on IBM PowerHA SystemMirror for i, refer to the following URL: <http://www.ibm.com/systems/power/software/availability/i/> .

IBM PowerHA SystemMirror for i now supports the copy services interfaces on the SAN Volume Controller (SVC) and Storwize storage servers, as well as the DS8000. IBM Lab Services provides additional support for many storage servers with their Full System FlashCopy Toolkit, as well as Advanced Copy Services for PowerHA on IBM i. For a current list of supported devices and technology, please refer to:

- PowerHA SystemMirror on i wiki – [www.ibm.com/developerworks/ibmi/ha](http://www.ibm.com/developerworks/ibmi/ha)
- Advanced Copy Services wiki - <https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20i%20Advanced%20Copy%20Services>

## ***12.2 DS4000 and DS5000***

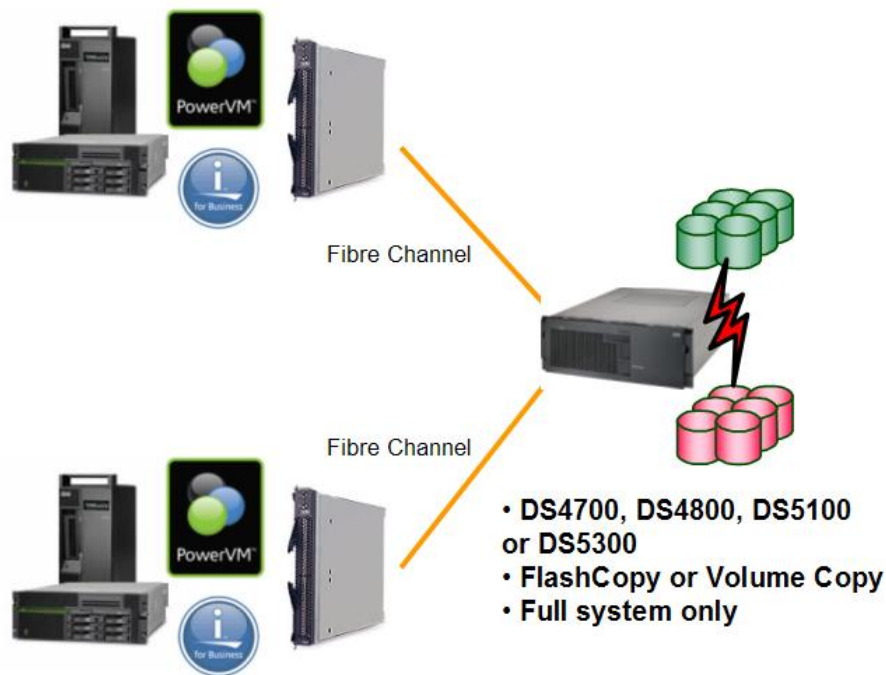
IBM has conducted some basic functional testing of DS4000 and DS5000 and Copy Services with IBM i as client of VIOS. In this section, you will find information on the scenarios tested and the resulting statements of support for using DS4000 and DS5000 Copy Services with IBM i.

### **12.2.1 FlashCopy and VolumeCopy**

#### **12.2.1.1 Test scenario**

Figure 11 shows the test environment used for FlashCopy and VolumeCopy:

- Production LPAR on IBM Power server or Power Blade
- IBM i as client of VIOS



- Backup LPAR on IBM Power server or Power Blade
- IBM i as client of VIOS

Figure 11 Test environment for Flashcopy and volume copy.

### 12.2.1.2 FlashCopy and VolumeCopy support statements

The use of DS4000 and DS5000 FlashCopy and VolumeCopy with IBM i as a client of VIOS is supported as outlined in this section. Note that to implement and use this solution, **multiple manual steps on the DS4000 or DS5000 storage subsystem, in VIOS and in IBM i are required**. The Lab Services Full System FlashCopy toolkit does offer support for DS4000 and DS5000 FlashCopy, although manual steps are still required. For current support information see the link within the [Copy Services General Support Information](#) section. There is no support for DS4000 or DS5000 FlashCopy within the IBM PowerHA® SystemMirror for i product. The components of the solution – DS4000 or DS5000 FlashCopy/VolumeCopy, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution will be provided by multiple IBM support organizations and not solely by the IBM i Support Center.

Support statements:

- DS4000 and DS5000 FlashCopy and VolumeCopy are supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Full-system FlashCopy and VolumeCopy when the production IBM i logical partition (LPAR) is powered off are supported .
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is in restricted state are supported.

- The DS4000 and DS5000 'disable' and 're-create' functions with full-system FlashCopy and VolumeCopy when the production IBM i LPAR is powered off or is in restricted state are supported.
- Full-system FlashCopy and VolumeCopy of the production IBM i logical partition (LPAR) after only using the IBM i 6.1 memory flush to disk (quiesce) function are **not supported**
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is running are **not supported**.
- FlashCopy and VolumeCopy of independent auxiliary storage pools (IASPs) are **not supported**.
- Having the production and backup IBM i LPAR under the same VIOS is **not supported**.

For assistance with using DS4000 and DS5000 FlashCopy and VolumeCopy with IBM i, contact IBM Systems Lab Services and Training using this website:  
<http://www.ibm.com/systems/services/labservices/contact.html>.

## 12.2.2 Enhanced Remote Mirroring (ERM)

### 12.2.2.1 Test scenario

Figure 12 shows the test environment used for Enhanced Remote Mirroring.

- **Production LPAR on IBM Power server or Power Blade**
- **IBM i as client of VIOS**



- **Backup LPAR on IBM Power server or Power Blade**
- **IBM i as client of VIOS**

Figure 12: the test environment used for ERM.

### 12.2.3 ERM support statements

The use of DS4000 and DS5000 Enhanced Remote Mirroring with IBM i as a client of VIOS is supported as outlined in this section. Note that to implement and use this solution, **multiple manual steps on the DS4000 or DS5000 storage subsystem, in VIOS and in IBM i are required**. Currently, no toolkit that automates this solution exists and it is not part of IBM PowerHA for IBM i. The components of the solution – DS4000 or DS5000 ERM, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution will be provided by multiple IBM support organization and not solely by the IBM i Support Center.

Support statements for ERM:

- DS4000 and DS5000 ERM is supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Synchronous ERM (DS4000 and DS5000 **Metro Mirror**) is supported.
- Asynchronous ERM with Write Consistency Groups (DS4000 and DS5000 **Global Mirror**) is supported.
- Asynchronous ERM (DS4000 and DS5000 Global Copy) is **not supported**.
- Full-system ERM (Metro Mirror and Global Mirror) for a planned switchover (IBM i production LPAR is powered off) is supported.
- Full-system ERM (Metro Mirror and Global Mirror) for an unplanned failover (IBM i production LPAR is running) is supported.
- Replication of IASPs is **not supported**.
- Suspending replication and accessing the replicated LUNs on backup site is **not supported**.

For assistance with using DS4000 and DS5000 Enhanced Remote Mirroring with IBM i, contact IBM Lab Services using the following website:  
<http://www.ibm.com/systems/services/labservices/contact.html>.

### 12.3 SAN Volume Controller (SVC) and Storwize Copy Services and IBM i

The SAN Volume Controller (SVC) and the Storwize storage servers (currently V7000, V5000, and V3700) all have the same copy services interfaces and all have the same copy services support statements.

The SVC and Storwize models are all supported through PowerHA SystemMirror for i as well as the Lab Services Full System FlashCopy toolkit. For current support statements for those products, please refer to the wikis referenced in the section above.


IBM has conducted some basic functional testing of SVC Copy Services with IBM i as client of VIOS. This testing was done without PowerHA or the Lab Services Full System FlashCopy toolkit. In this section, you will find information on the scenarios tested and the resulting statements of support for using SVC Copy Services with IBM i.

#### 12.3.1 FlashCopy

##### 12.3.1.1 Test scenario

Figure 13 shows the test environment used for FlashCopy.

- Production LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS



- Backup LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS

Figure 13: The test environment used for FlashCopy.

### 12.3.1.2 FlashCopy support statements

The use of SVC and Storwize FlashCopy with IBM i as a client of VIOS is supported as outlined below. **Please note** that to implement and use this solution, **multiple manual steps in SVC, in VIOS and in IBM i are required**. Both the Lab Services Full System FlashCopy toolkit and PowerHA SystemMirror for i provide support for SVC and Storwize FlashCopy, minimizing the steps required. If not using one of those solutions, the components of the solution – SVC FlashCopy, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for this solution outside of the Lab Services toolkit or PowerHA will be provided by multiple IBM support organizations and not solely by the IBM i Support Center.

Support statements for SVC and Storwize FlashCopy:

- FlashCopy is supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Full-system FlashCopy when the production IBM i logical partition (LPAR) is powered off is supported .
- Full-system FlashCopy when the production IBM i LPAR is in restricted state is supported.
- Full-system FlashCopy of the production IBM i logical partition (LPAR) after using the IBM i 6.1 memory flush to disk (quiesce) function is supported.
- Full-system FlashCopy when the production IBM i LPAR is running is **supported, but not recommended as much recovery may be needed on the FlashCopy target**.
- FlashCopy between VDisks using heterogeneous storage subsystems is supported.

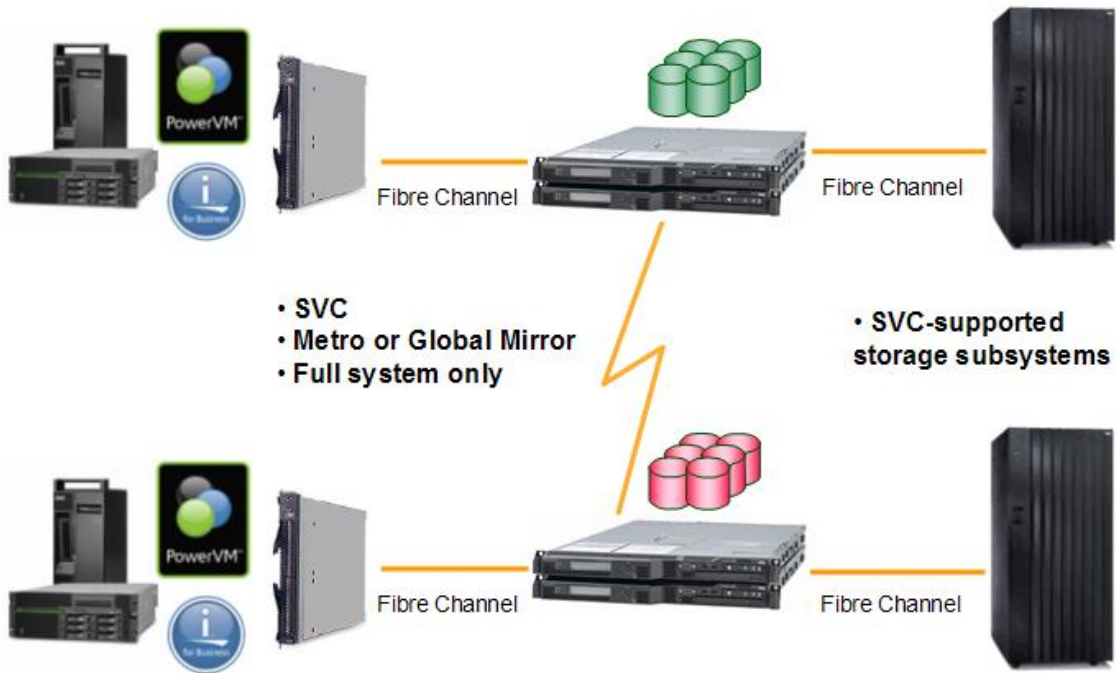
For assistance with using SVC FlashCopy with IBM i, contact IBM Lab Services or refer to the following website: <http://www.ibm.com/systems/services/labservices/contact.html>.

## 12.3.2 Metro and Global Mirror

### 12.3.2.1 Test scenario

Figure 14 shows the test environment used for Metro and Global Mirror.

- Production LPAR on IBM Power server or IBM BladeCenter JS12/JS22\*
- IBM i as client of VIOS



- Backup LPAR on IBM Power server or IBM BladeCenter JS12/JS22
- IBM i as client of VIOS

Figure 14: The test environment used for Metro Mirror and Global Mirror.

### 12.3.2.2 Metro Mirror and Global Mirror support statements

The use of SVC and StorwizeMetro Mirror and Global Mirror with IBM i as a client of VIOS is supported as outlined in the following section. Note that to implement and use this solution, **multiple manual steps in SVC or Storwize, in VIOS and in IBM i are required**. PowerHA SystemMirror for i provides support for SVC and Storwize Metro Mirror and Global Mirror, minimizing the steps required. If not using the PowerHA solutions, the components of the solution – SVC or Storwize Metro or Global Mirror, VIOS and IBM i – must be managed separately and require the corresponding skill set. Also note that support for a solution outside of PowerHA will be provided by multiple IBM support organization and not solely by the IBM i Support Center.

Support statements for SVC and Storwize Metro and Global Mirror:

- Both Metro Mirror and Global Mirror are supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Both Metro Mirror and Global Mirror are supported for a planned switchover (IBM i production partition is powered off).

- Both Metro Mirror and Global Mirror are supported for an unplanned failover (IBM i production partition is running).
- Suspend and resume in case of link failure between SVC clusters is supported for both Metro Mirror and Global Mirror.

For assistance with using SVC or Storwize Metro and Global Mirror with IBM i, contact IBM Lab Services using the following website:

<http://www.ibm.com/systems/services/labservices/contact.html>.

## **12.4 XIV Copy Services and IBM i**

IBM has conducted some basic functional testing of XIV System Storage Copy Services with IBM i as client of VIOS. Full system FlashCopy of XIV is supported via the IBM Lab Services Full System FlashCopy Toolkit. For details, see the wiki link in Copy Services [General Support Information](#).

The following lists are the resulting statements of support for using XIV Copy Services with IBM i.

Snapshots and Volume Copy:

- XIV snapshots and Volume Copy are supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades.
- Full-system snapshots and Volume Copy when the production IBM i logical partition (LPAR) is powered off are supported.
- Full-system snapshots and Volume Copy when the production IBM i LPAR is in restricted state are supported.
- Full-system snapshots and Volume Copy of the production IBM i logical partition (LPAR) after using the IBM i 6.1 memory flush to disk (quiesce) function are supported.
- Full-system snapshots and Volume Copy when the production IBM i LPAR is running are **not supported**.
- Snapshots and Volume Copy of iASPs are **not supported**.

Remote Mirror:

- XIV Remote Mirror (synchronous replication) is supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades
- XIV asynchronous replication is also supported by IBM i as a client of VIOS on both IBM Power servers and IBM Power blades
- Only full-system replication is supported
- Replication of iASPs with Remote Mirror is **not supported**
- Remote Mirror is supported for a planned switchover (IBM i production partition is powered off)
- Remote Mirror is supported for an unplanned failover (IBM i production partition is running)

## **12.5 DS5000 Direct attach Copy Services**

IBM has conducted some basic functional testing of DS5100 and DS5300 Copy Services when directly attached to IBM i. This section summarizes the resulting support statements. Note that using Copy Services when directly attaching DS5100 and DS5300 storage systems involves manual steps in both the DS Storage Manager GUI and IBM i. PowerHA for IBM i does not support DS5100 and DS5300 Copy Services when directly attached to IBM i (or through VIOS).

FlashCopy and VolumeCopy:

- FlashCopy and VolumeCopy are supported when directly attaching DS5100 and DS5300 to IBM i on POWER6 or higher processor-based Power servers.
- As mentioned earlier, Power blades are **not supported** for direct IBM i attachment to DS5100 and DS5300 in general.

- Full-system FlashCopy and VolumeCopy when the production IBM i logical partition (LPAR) is powered off are supported.
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is in restricted state are supported.
- Full-system FlashCopy and VolumeCopy of the production IBM i logical partition (LPAR) after using the IBM i 6.1 memory flush to disk (quiesce) function are **not supported**.
- Full-system FlashCopy and VolumeCopy when the production IBM i LPAR is running are **not supported**.
- FlashCopy and VolumeCopy of Independent Auxiliary Storage Pools (IASPs) are supported.

Enhanced Remote Mirroring (ERM):

- DS5100 and DS5300 ERM with direct attachment to IBM i is supported on IBM Power servers (but **not** on IBM Power blades).
- Synchronous ERM (DS5100 and DS5300 **Metro Mirror**) is supported.
- Asynchronous ERM with Write Consistency Groups (DS5100 and DS5300 **Global Mirror**) is supported.
- Asynchronous ERM (DS5100 and DS5300 **Global Copy**) is **not supported**.
- Full-system ERM (Metro and Global Mirror) for a planned switchover (IBM i production LPAR is powered off) is supported.
- Full-system ERM (Metro and Global Mirror) for an unplanned failover (IBM i production LPAR is running) is supported.
- ERM (Metro and Global Mirror) with IASPs for a planned switchover (IBM i production LPAR is powered off) is supported.
- ERM (Metro and Global Mirror) with IASPs for an unplanned failover (IBM i production LPAR is running) is supported.

## ***12.6 NPIV attached storage Copy Services***

NPIV support is now available for the DS8000, SVC, and Storwize storage servers.

The main differences between NPIV and VSCSI I/O virtualization for IBM i are:

- With NPIV, storage and tape devices report into IBM i with their specific characteristics, as opposed to generic SCSI devices.
- IBM i has also gained NPIV access to storage and tape devices to which it could already attached directly, benefiting from transparent FC adapter sharing.

As a result, both the storage subsystem and LUN types involved with NPIV for IBM i are identical to those already supported by IBM PowerHA System Mirror for IBM i. PowerHA continues to manage both Copy Services on the storage server as well as the clustering interfaces in IBM i from a single interface. Therefore, DS8000, SVC and Storwize storage system Copy Services are supported by PowerHA for IBM i when using NPIV, including FlashCopy, Metro Mirror and Global Mirror.



## 13 Appendix:

### 13.1 Additional resources

These websites provide useful references to supplement the information contained in this paper

#### 13.2 IBM i

- IBM i on a Power Blade Read-me First:  
<http://www.ibm.com/systems/power/hardware/blades/ibmi.html>.
- IBM STG Lab Services: <http://www.ibm.com/systems/services/labservices/contact.html>.
- Logical Partitioning Guide:  
<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphath.pdf>.
- IBM i installation:  
[http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahc/rzahc1.htm&tocNode=int\\_216451](http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahc/rzahc1.htm&tocNode=int_216451)
- IBM i PTF installation:  
<http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzam8/rzam8fixconcepts1.htm>
- Latest recommended IBM i PTFs: [http://www-912.ibm.com/s\\_dir/slkbase.nsf/recommendedfixes](http://www-912.ibm.com/s_dir/slkbase.nsf/recommendedfixes)
- IBM Fix Central: <http://www-912.ibm.com/eserver/support/fixes/>
- IBM Systems Workload Estimator: <http://www-304.ibm.com/systems/support/tools/estimator/index.html>
- IBM i networking:  
<http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzajy/rzajyoverview.htm>
- ECS over LAN:  
[http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzaji/rzaji\\_setup.htm](http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzaji/rzaji_setup.htm)
- IBM i virtual client partitions topics in the IBM i Information Center:  
<http://publib.boulder.ibm.com/infocenter/series/v7r1m0/index.jsp?topic=/rzahc/rzahcbladei5limits.htm>
- IBM PowerHA for IBM i: <http://www-304.ibm.com/jct03001c/systems/power/software/availability/i5os.html>.

#### 13.3 Storage

- *IBM System Storage DS4000 and Storage Manager V10.30* (Redbook):  
<http://www.redbooks.ibm.com/redpieces/abstracts/sg247010.html?Open>
- DS4000 Website: <http://www.ibm.com/systems/storage/disk/ds4000/index.html>
- *Performance Capabilities Reference* manual (chapter 6 for best practices for IBM i and VIOS): <http://www.ibm.com/systems/i/solutions/perfmgmt/resource.html>
- SDD-PCM driver:  
[http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D430&uid=ssg1S4000201&loc=en\\_US&cs=utf-8&lang=en](http://www.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D430&uid=ssg1S4000201&loc=en_US&cs=utf-8&lang=en)
- IBM Redbooks site: <http://www.redbooks.ibm.com>
- *IBM i and IBM System Storage: A Guide to Implementing External Disks on IBM i* (Redbook): <http://www.redbooks.ibm.com/abstracts/sg247120.html?Open>.
- *Implementing IBM Tape in i5/OS* (Redbook):  
<http://www.redbooks.ibm.com/abstracts/sg247440.html?Open>.
- *IBM Midrange System Storage Hardware Guide*(Redbook):  
<http://www.redbooks.ibm.com/abstracts/sg247676.html?Open>.

- [IBM System Storage DS3500: Introduction and Implementation Guide:](http://www.redbooks.ibm.com/redpieces/pdfs/sq247914.pdf)  
<http://www.redbooks.ibm.com/redpieces/pdfs/sq247914.pdf>

## 13.4 VIOS

- *PowerVM Editions Guide:*  
<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/arecu/arecukickoff.htm>
- *Advanced POWER Virtualization on IBM System p5: Introduction and Configuration* (Redbook): <http://www.redbooks.ibm.com/abstracts/sq247940.html?Open>
- VIOS command reference:  
[http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1\\_vios\\_commandslist.htm](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphb1/iphb1_vios_commandslist.htm)
- VIOS Datasheet:  
<http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/datasheet.html>
- *PowerVM Managing and Monitoring* (Redbook):  
<http://www.redbooks.ibm.com/abstracts/sq247590.html?Open>.

## 13.5 SVC

- SVC overview Website: <http://www.ibm.com/systems/storage/software/virtualization/svc>
- Interoperability information:  
<http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html>
- *Implementing the IBM System Storage SAN Volume Controller V5.1* (Redbook):  
<http://www.redbooks.ibm.com/abstracts/sq246423.html?Open>.
- *SAN Volume Controller Best Practices and Performance Guidelines* (Redbook):  
<http://www.redbooks.ibm.com/abstracts/sq247521.html?Open>.
- *SVC V4.3.0 Advanced Copy Services* (Redbook):  
<http://www.redbooks.ibm.com/redpieces/abstracts/sq247574.html?Open>.
- *IBM SAN Volume Controller 4.2.1 Cache Partitioning* (Redpiece):  
<http://www.redbooks.ibm.com/abstracts/redp4426.html?Open>.

## 14 Trademarks and disclaimers

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Refer to your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L (logo), AIX 6 (logo), alphaWorks, AS/400, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DirectTalk, Domino, DYNIX, DYNIX/ptx, e business (logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, IBM i, IBM i (logo), IBM, IBM (logo), ibm.com, IBM Business Partner (logo), Informix, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, OpenPower, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, Passport Advantage, POWERparallel, Power PC 603, Power PC 604, PowerPC, PowerPC (logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, Quick Place, Rational, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Sametime, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SpaceBall, System/390, The Engines of e-business, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, Tivoli Ready (logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Advanced Micro-Partitioning, AIX 5L, AIX PVMe, AS/400e, Calibrated Vectors Cooling, Chipopper, Chipkill, Cloudscape, DataPower, DB2 OLAP Server, DB2 Universal Database,

DFDSM, DFSORT, DS4000, DS6000, DS8000, e-business (logo), e-business on demand, EnergyScale, Enterprise Workload Manager, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, General Purpose File System, GigaProcessor, GPFS, HACMP, HACMP/6000, IBM Systems Director Active Energy Manager, IBM TotalStorage Proven, IBMLink, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, On Demand Business logo, POWER, PowerExecutive, PowerVM, PowerVM (logo), Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power PC, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, pure XML, Quickr, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, System i, System i5, System p, System p5, System Storage, System z, System z9, S/390 Parallel Enterprise Server, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, Workload Partitions Manager, X-Architecture, z/Architecture, z/9.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>. The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries or both.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

Intel, Itanium, Pentium are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

AltiVec is a trademark of Freescale Semiconductor, Inc.

Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Other company, product and service names may be trademarks or service marks of others.