Tivoli Storage Manager FastBack for Bare Machine
Recovery
Version 6.1.7

User's Guide

IBM

Tivoli Storage Manager FastBack for Bare Machine
Recovery
Version 6.1.7

User's Guide

IBM

> **Note**
>
> Before using this information and the product it supports, read the information in "Notices" on page 37.

Contents

Figures

Tables

Preface

This publication helps you install and use Tivoli Storage Manager FastBack for Bare Machine Recovery.

Who should read this guide

This publication provides instructions for a user to install, configure, and use Tivoli Storage Manager FastBack for Bare Machine Recovery.

Publications

Tivoli Storage Manager FastBack publications and other related publications are available online.

You can search publications in the Tivoli Storage Manager FastBack Information Center: http://publib.boulder.ibm.com/infocenter/tsmfbinf/v6/index.jsp

You can download PDF versions of publications from the Tivoli Storage Manager FastBack Information Center or from the IBM Publications Center at http://www.ibm.com/shop/publications/order/.

Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: http://www.ibm.com/support/entry/portal/. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli® technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and to interact with others who use IBM® storage products.

Tivoli software training and certification
Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at http://www.ibm.com/software/tivoli/education/.

Tivoli Support Technical Exchange
Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community
Interact with others who use IBM storage management products at http://www.ibm.com/developerworks/servicemanagement/sm/index.html.

Global Tivoli User Community
Share information and learn from other Tivoli users throughout the world at http://www.tivoli-ug.org/.

IBM Education Assistant
View short "how to" recordings designed to help you use IBM software products more effectively at http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp.

Searching knowledge bases

If you have a problem you can search for information in a knowledge base.

Search the Tivoli Storage Manager FastBack Information Center at http://publib.boulder.ibm.com/infocenter/tsmfbinf/v6/index.jsp.

Search the internet

If you cannot find an answer to your question in the information center, search the Internet for the latest, most complete information that might help you resolve your problem.

To search multiple Internet resources for your product, go to the Tivoli Support portal at https://www.ibm.com/support/entry/portal/overview/software/tivoli/tivoli_brand_support_%28general%29 and search support for the product. From this section, you can search a variety of resources including:
- IBM technotes
- IBM downloads
- IBM Redbooks®
- Forums and newsgroups

Using IBM Support Assistant

At no additional cost, you can install on any workstation the IBM Support Assistant, a stand-alone application. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use.

The IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:
- Support links
- Education links
- Ability to submit problem management reports

For more information, see the IBM Support Assistant Web site at http://www.ibm.com/software/support/isa/.

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at http://www.ibm.com/support/docview.wss?uid=swg27012689.

Finding product fixes

A product fix to resolve your problem might be available from the IBM Support Assistant website.

About this task

To check what fixes are available for your product, follow these steps:

Procedure
- From the IBM Support Assistant Web site at http://www.ibm.com/support/entry/portal/, click **Downloads**.
- Click **Search for recommended fixes**.
- Choose content filters to find fixes for your product level and operating system.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

About this task

To sign up to receive notifications about IBM products, follow these steps:

Procedure
1. From the IBM Support Assistant Web site at http://www.ibm.com/support/entry/portal/, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.
6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract and if you are authorized to submit problems to IBM.

About this task

To obtain help from IBM Software Support, complete the following steps:

Procedure
1. Ensure that you have completed the following prerequisites:
   a. Set up a subscription and support contract.
   b. Determine the business impact of your problem.
   c. Describe your problem and gather background information.
2. Follow the instructions in "Submitting the problem to IBM Software Support" on page xiii.

Setting up a software maintenance contract

Set up a software maintenance contract. The type of contract that you need depends on the type of product you have.

Procedure

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as IBM DB2® and IBM WebSphere® products that run on Microsoft Windows or UNIX operating systems), enroll in IBM Passport Advantage® in one of the following ways:
  - **Online:** Go to the Passport Advantage Web page at http://www.ibm.com/software/lotus/passportadvantage/, click **How to enroll**, and follow the instructions.
  - **By Phone:** For the phone number to call in your country, go to the IBM Software Support Handbook Web page at http://techsupport.services.ibm.com/guides/contacts.html and click **Contacts**.
- For server software products, you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for server software products, go to the IBM Technical support advantage Web page at http://www.ibm.com/servers/eserver/techsupport.html.

What to do next

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. For a list of telephone numbers of people who provide support for your location, go to the Software Support Handbook page at http://www.ibm.com/support/customercare/sas/f/handbook/home.html.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Severity 1	**Critical** business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	**Significant** business impact: The program is usable but is severely limited.
Severity 3	**Some** business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	**Minimal** business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.

- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online
> Go to the IBM Software Support website at http://www.ibm.com/ support/entry/portal/Open_service_request/Software/ Software_support_(general). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone
> For the telephone number to call in your country, go to the IBM Software Support Handbook at http://www14.software.ibm.com/webapp/set2/sas/ f/handbook/home.html and click **Contacts**.

Chapter 1. Overview

Tivoli Storage Manager FastBack for Bare Machine Recovery is part of the Tivoli Storage Manager FastBack product family.

Tivoli Storage Manager FastBack for Bare Machine Recovery is a solution for Windows operating systems that you can use to recover entire systems to a comparable server, to a new server with different hardware, or to a virtual machine. You can also use Tivoli Storage Manager FastBack for Bare Machine Recovery to migrate a server.

Tivoli Storage Manager FastBack supports Tivoli Storage Manager FastBack for Bare Machine Recovery on both 32 and 64-bit operating systems and processors.

A bare machine recovery has the following advantages:
- Reduce the amount of time it takes for you to get to the point where you start restoring data.
- Automates some of the steps, for example, re-partitioning your hard drives.

One way to reduce the amount of downtime after a system failure is to avoid installing the operating system on your hard drive. Tivoli Storage Manager FastBack for Bare Machine Recovery helps you to restore your operating system by booting an operating system from a CD into the computer memory, as opposed to booting the operating system from a hard disk. Booting the operating system from a hard disk is not feasible because you are recovering a system. When you run Tivoli Storage Manager FastBack for Bare Machine Recovery, a boot CD, referred to in this document as the FastBack for Bare Machine Recovery CD, is required.

Cloned disks for deployment can also be created. In this case, some manual changes in the system names and, perhaps, Active Directory definitions might be required.

The following situations require the use of Tivoli Storage Manager FastBack for Bare Machine Recovery:
- Recovery from a corrupted operating system.
- Recovery from failed system hardware.
- Restoring a disk to the same or a different disk on the same computer due to disk failure.
- Creating clone disks (including the operating system) for deployment.
- Restoring the operating system and environment due to wrong installations or patches (for example, restore to the same disk and same boot volume).

In a cluster environment, when running bare machine recovery, direct bare machine recovery to a clustered disk, including a quorum disk, is not supported. However, you can remove the disk from the cluster before starting a bare machine recovery. After bare machine recovery is complete, you can return the disk to the cluster.

Supported environments

Before using Tivoli Storage Manager FastBack for Bare Machine Recovery, ensure to meet the processor and operating system requirements.

Tivoli Storage Manager FastBack for Bare Machine Recovery supports the following environments:

- x86 processors (32 bit) with Microsoft Windows 32-bit operating systems
- x64 (AMD64 and EM64T) processors with Microsoft Windows 32-bit or 64-bit operating systems
- UEFI platform machines (IA64 - Intel Itanium processors, or X86 processors)

Support for virtual machines and virtualization

When you run Tivoli Storage Manager FastBack software on virtual machines and through virtualization, the guest system must run an operating system supported by the Tivoli Storage Manager FastBack component. For example, the guest system where the FastBack Server is installed must run one of the supported operating systems for FastBack Server.

The following list summarizes support for virtual machines and virtualization:

VMware ESX guest
> Tivoli Storage Manager FastBack products and components are supported for backup and recovery within the VMware ESX virtual guest.

Microsoft Hyper-V virtual guest
> Tivoli Storage Manager FastBack products and components are supported for backup and recovery within the Microsoft Hyper-V virtual guest. Backups of the Hyper-V virtual machines from the parent partition by using Microsoft' Volume Shadow Copy Services (VSS) is not supported.
>
> If you complete a bare machine recovery to a Hyper-V virtual machine, use the following list of tips:
>
> - The target virtual machine for the bare machine recovery must have an established network adapter. You can add a legacy network adapter to the virtual machine by going to the Hyper-V Manager, clicking **Settings** for the virtual machine, and then clicking **Add Hardware** > **Legacy Network Adapter**.
> - To have the mouse available on the Hyper-V virtual machine, install *Hyper-V Manager*. Hyper-V Manager requires the Microsoft Windows Vista, Service Pack 1, or Microsoft Windows Server 2008.
> - The boot partition must be on a virtual IDE disk. Hyper-V cannot boot from a SCSI disk.

x86 processors (32-bit) with Windows 32-bit operating systems

The following list describes what you need to complete the disk restore processing by using Tivoli Storage Manager FastBack for Bare Machine Recovery:

- The user credentials entered for disk restoration must have privileges for Tivoli Storage Manager FastBack for Bare Machine Recovery on the volumes and disks to be restored.
- The FastBack for Bare Machine Recovery CD requires a minimum of 512 MB of RAM.

- The capacity of the destination disk must support the restore process operation. You need the same size or larger. If additional space is available after the restore process, it can be configured manually for use after the restore process.
- After a successful Tivoli Storage Manager FastBack for Bare Machine Recovery of a system disk, Windows can require product activation. For Windows XP and later, the Windows XP activation feature requires that the operating system on restored disks be reactivated because previous activation was relevant to a specific disk ID.

When restoring a disk with Tivoli Storage Manager FastBack for Bare Machine Recovery, you can complete the following tasks:
- Restore any successful snapshot, both full and incremental.
- Boot your system, connect to the network, and run Tivoli Storage Manager FastBack for Bare Machine Recovery that connects to a FastBack Server.
- When a disk restore process is initiated, create a master boot record with up to three primary partitions followed by logical partitions. The original layout of the disk is not maintained.
- Take snapshots that can be volumes of dynamic or basic disks. However, the restore process creates a basic disk.
- Create any combination of user-selected volumes. Retain the original disk-layout to avoid unpredictable operating system behavior.
- If the restore of one or more volumes fails, the empty or allocated partition is available on the disk and it can be allocated manually by the Microsoft Disk Manager tool.

If the destination disk is dynamic or has volumes, all previous data on the disk is erased. The restored disk is not identified by the server that owns the destination disk until a refresh is complete. To refresh, go to the following menu: **MMC** > **Storage** > **Disk Management** > **Rescan Disks**.

Test the FastBack for Bare Machine Recovery CD in your environment during regular operation (before a disaster occurs) to confirm that all the necessary hardware drivers are already on the CD. Tivoli Storage Manager FastBack for Bare Machine Recovery of EISA partitions must be included in the backup snapshot in order for the system to be properly restored.

x64 (AMD64 and EM64T) processors with Microsoft Windows 32-bit or 64-bit operating systems

The procedure for using Tivoli Storage Manager FastBack for Bare Machine Recovery on the x64 (AMD64 and EM64T) processors with Windows 32-bit or 64-bit operating systems is identical to using Tivoli Storage Manager FastBack for Bare Machine Recovery on the x32 processors with Windows 32-bit operating systems.

UEFI platform machines (IA64 Intel Itanium)

UEFI systems depend on NVRAM parameters and a different disk layout (GUID partition table disk) than the layout used on MBR disk BIOS systems. Tivoli Storage Manager FastBack for Bare Machine Recovery on an UEFI system requires that you configure the NVRAM, EFI partition, utility partitions, and system partition for the UEFI system, and restore the original system data partition. Restoring the original system data partition requires that the boot disk that is being restored is accessible by a second, functioning UEFI based system that runs the FastBack Client software.

To configure the NVRAM, EFI partition, utility partitions, and system partition for the UEFI system, complete the following steps:

1. Identify the original volume that you want to restore. For example, restore the `C:\` volume on the UEFI system. In the following steps, the volume you want to restore is referenced as the original system. The EFI partition does not need to be backed up. The EFI partition is created later.

2. Physically install the target hard disk on the original system. This hard disk must be large enough to contain the EFI partition, utility partitions, and system partitions of the original system. The target disk must be at least the same size as the source disk. The disk must be installed as a primary disk on the UEFI system. Because the source disk is faulty, the source disk is to be removed from the UEFI system.

3. Boot the original system with the utility CD from the manufacturer. The disk is used for installing a new UEFI system.

4. Run an express setup, or an equivalent type of setup.

5. Use the following procedure to install a regular Windows Itanium or 2008 version on this hard disk:

   a. After the Windows setup copies the setup files, allow the installation process continue to build the correct partition.

   b. The Windows installation process prompts you to choose a type of installation: `Express` or `Manual`. Choose **Manual**. This installation choice is required so you can later specify partition size settings.

   c. When the installation process asks for the size of the Windows partition, make the partition at least the same size as the Windows partition on the original system.

   d. Complete the installation process.

6. Shut down the original system and remove the hard disk. You need this hard disk when restoring the original system data partition.

To restore the original system data partition, complete the following steps:

1. Install the hard disk that you removed at the end of the configuring the NVRAM, EFI partition, utility partitions, and system partition for the UEFI system section. This hard disk is a second disk, on a different UEFI system, that is running the FastBack Client. In the following steps, the second UEFI system is referenced as the functioning system.

   To install this hard disk as a second disk on the functioning system, the functioning system must be shut down before the hard disk installation process. In addition, restart the functioning system after the hard disk installation process is complete.

2. Configure the FastBack Client on the functioning system to connect to the server that contains the system partition from the original system. This system partition is the partition that you want to restore.

3. Run a volume restore of the system volume on the second hard disk in the functioning system from the FastBack Manager snapshot monitor window. Restore the snapshot on the system partition. Do not restore to another partition on the disk. You are going to overwrite the default system volume installed by Windows.

4. After the volume restore process, shut down the functioning system.

5. Remove the second disk.

6. Reinstall the hard disk as the primary disk in the original system.

7. Restart the original system.

8. Connect the FastBack Client on the original system to the appropriate server.
9. Restore additional partitions on the disk by using the standard volume restore process from FastBack Manager.

Note: The FastBack for Bare Machine Recovery CD does not load on UEFI systems.

UEFI platform machine (X86 processors)

On X86_64 UEFI machines, there is no need to physically remove the hard disk. After installing the operating system to configure the partitions, complete the following steps:

1. Start the computer with the FastBack for Bare Machine Recovery CD.
2. Choose the drive snapshot (usually volume C:) that you want to restore.
3. Use the advanced (volume restore) option of Tivoli Storage Manager FastBack for Bare Machine Recovery to restore the system volume to the system volume partition created by the windows installation. For more information, see step 9 of the "FastBack for Bare Machine Recovery CD for bare machine recovery on your Windows system" section).
4. After the volume restore is complete remove the FastBack for Bare Machine Recovery CD and restart the computer.

Note: On UEFI virtual machines, when installing the Windows operating system the computer must be configured to UEFI boot. Mark the machine as UEFI in the configuration settings under boot option. When loading the FastBack for Bare Machine Recovery CD configure the virtual machine to BIOS boot (not UEFI boot). When the restore process is over, configure the machine back to UEFI boot.

Access permissions

Tivoli Storage Manager FastBack for Bare Machine Recovery can be used from sources and to destinations that the currently logged user is authorized to access.

`SuperAdmin` users have unlimited restore access. These users can be either Active Directory users, or FastBack Manager domain users that belong to the `SuperAdmin` group.

Active Directory users that are not `SuperAdmin` users require Share and ACL read permissions to all source volumes root directories in the disk, and Share and ACL modify permissions to all destination volumes root directories in the disk. The source permissions are those that were granted during backup, while the destination permissions are those that were granted at restore time. Share permissions are not required if both Source and Destination are on the same computer as the Tivoli Storage Manager FastBack, from which the restore operation is being initiated.

Note: Permissions are not checked for target disks without partitions. Only `Admin` users can access target disks with unformatted partitions, and source and target volumes that are not mapped to any Clients.

Figure 1 on page 6 display the access permissions.
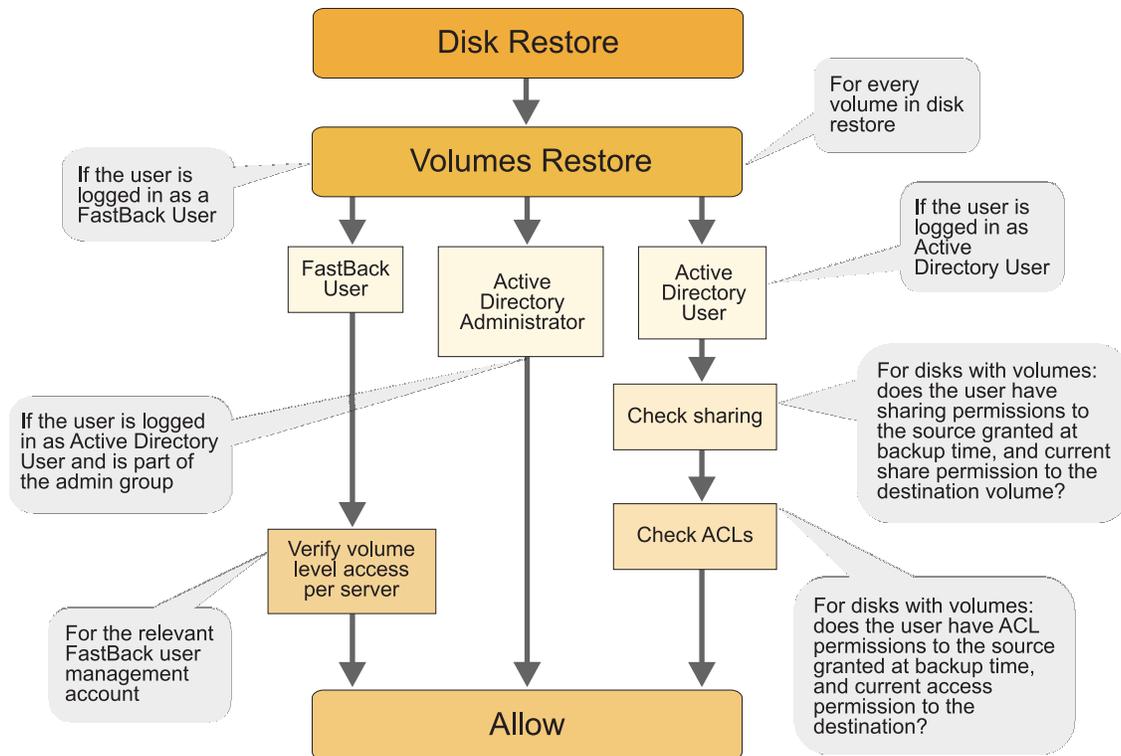
Figure 1. Tivoli Storage Manager FastBack for Bare Machine Recovery access permissions

Using the FastBack for Bare Machine Recovery CD

To restore a volume to a computer in situations where the disk or the operating system was corrupted, or in other situations that caused server failure you must create a FastBack for Bare Machine Recovery CD to use with Tivoli Storage Manager FastBack for Bare Machine Recovery.

The procedure consists of two phases:
- If necessary, repairing or replacing the damaged hardware.
- Restore with Tivoli Storage Manager FastBack for Bare Machine Recovery on the FastBack for Bare Machine Recovery CD. You must create the FastBack for Bare Machine Recovery CD.

When creating the FastBack for Bare Machine Recovery CD, consider the following items:
- Each user must create their own customized CD that includes storage drivers.
- When using the FastBack for Bare Machine Recovery CD, only Active Directory group users and Tivoli Storage Manager FastBack administrators can restore.
- The FastBack Server service does not need to run to use Tivoli Storage Manager FastBack for Bare Machine Recovery.
- If the BIOS on the booting computer is set to boot from a drive other than the DVD drive, change the BIOS to boot from the DVD drive. Then reinsert the CD.
- You can create the FastBack for Bare Machine Recovery CD on any operating system that is supported by the Tivoli Storage Manager FastBack Client.

Any-to-Any HW Restore utility

If the restored data structure does not match the hardware configuration of the destination disk, the operating system cannot run. In addition, after a restart, a blue screen is displayed on the restored system.

To correct this problem, use the `Any-to-Any HW Restore` utility, included on the FastBack for Bare Machine Recovery CD. This utility provides various diagnostic tools. You can use the utility to identify and correct problems related to data structure or mass storage driver differences. For information about how to use the `Any-to-Any HW Restore` utility, see "Using the `Any-to-Any Hardware Restore` utility" on page 16.

CAUTION:
Depending on the symptoms and troubleshooting progress, use the `Any-to-Any HW Restore` utility features with caution to solve the dissimilar hardware issues and complete the restore. Perform step-by-step changes by using the utility, while restarting after each individual change. For instance, install a driver and restart to see how it affects the system. Make another change and restart again.

Customizing volume layout

You can use Tivoli Storage Manager FastBack for Bare Machine Recovery to customize volume layout. Disks can also be restored for other purposes than system-level restore.

The following information lists the examples where you can restore disks for other purposes:
- Various volumes that are originally on different disks or different servers.
- Various volumes that are snapped at different times.
- Same volume at different points in time.

If the disk to be restored is a boot disk, unpredictable behavior of the operating system can occur under the following conditions:
- Your operating system originally exists in a specific volume, for example, the second volume on the disk. The operating system is restored to another volume, for example, the first volume on the disk. Change the `boot.ini` file or edit the boot configuration data file by using `BCDEDIT.exe` to reflect the changes.
- Your swap files originally exists on a number of volumes (for example, `C:` and `D:`), and only some of the volumes were restored (for example, `C:` only). Recover all volumes that have swap files.
- You originally had a multi-boot system (for example, three operating systems on the `C`, `D`, and `E` drives) and you restore only one volume. Change the `boot.ini` file or edit the boot configuration data file by using `BCDEDIT.exe` to reflect the changes. For more information about how to configure the ARC naming paths that are used in the `boot.ini` file, go reference the *Microsoft Knowledge Base Article 102873: BOOT.INI and ARC Path Naming Conventions and Usage*.
- Software applications that are originally distributed over a number of disks were partially restored.

Installing Tivoli Storage Manager FastBack for Bare Machine Recovery

About this task

To install IBM Tivoli Storage Manager FastBack for Bare Machine Recovery, complete the following steps:

Procedure

1. Download the code package.
2. Navigate to the folder where the code package is saved.
3. Start the installation process by running the setup.exe file.
4. Select the desired language.
5. The welcome page is displayed. Click **Next**.
6. The Software License Agreement page is displayed. Read the terms of the license agreement. To accept the license agreement, click **Yes**. You must accept the terms of the license agreement to continue the installation.
7. Select a setup type.
   a. For a `Complete` installation, select `Complete` and click **Next**.
   b. For a `Custom` installation, select `Custom` and click **Next**. Then select the components to be installed and the installation path and click **Next**.

      Note: If you chose to install the `Linux for BMR` component, the installation path should not be changed from the default path.
8. Click **Install**.

   Note: If you chose to install the `Linux for BMR` component, an error will be displayed if FastBack Mount is not installed.
9. Click **Finish**.

Chapter 2. Creating a FastBack for Bare Machine Recovery CD for your Windows system

To build the FastBack for Bare Machine Recovery CD you must download several software packages and complete customization procedures to fit the CD to your environment.

Before you begin

Before you can create a FastBack for Bare Machine Recovery CD for your Windows system, ensure you have the following software and hardware components:

- CD writing (burning) software.
- Recordable CD drive.
- Internet connection for downloading software packages from the Web.
- A system with a 32-bit Windows XP, Windows 2003, or Windows Vista operating system installed.
- 1 GB of free disk space to build the CD image.

About this task

To create a FastBack for Bare Machine Recovery CD, complete the following steps:

Procedure

1. Download the Tivoli Storage Manager FastBack for Bare Machine Recovery base code package that you downloaded from either Passport Advantage or the IBM Support Fix Central Web site.

2. Download and install the Microsoft Windows Automated Installation Kit (AIK) on a 32-bit operating system. You can download the latest Microsoft Windows AIK from the Microsoft Web site: http://www.microsoft.com/downloads/details.aspx?familyid=94BB6E34-D890-4932-81A5-5B50C657DE08&displaylang=en.

   After the download is complete, create a CD from the downloaded image using your CD burning software and install the AIK software following the instructions on the Windows Automated Installation Kit (AIK) Web site.

   In order to install additional languages, you need this Windows AIK CD that you just created. Microsoft PE is part of the AIK package you download and install.

3. Create a temporary work folder on a local disk. In this document, the example temporary work folder is labeled `\FastBackBMR` and is located at the root of `C:`. The work area is used to create the bootable CD image. You only need to create the parent folder for this step. The rest of the folders contained in the parent folder are created in the subsequent steps.

4. (Optional) Download approved storage drivers for Microsoft Windows PE. If you have systems with storage that are not supported by the Windows 2008 PE, you have to add your own drivers. For Microsoft PE to recognize your storage devices, you need to customize your CD image with storage drivers that are suitable for your environment.

To download approved storage drivers, see http://www.ibm.com/support/ docview.wss?rs=3508&context=SS9NW2&uid=swg21385896. Ensure that you download storage drivers to match both your existing and future systems.

Copy the drivers into the `\FastBackBMR\PEDrivers\StorageDrivers` folder by completing the following steps:

Note: In addition to the following list of steps, each storage driver might require additional steps to complete installation and configuration. For complete instructions, see instructions at http://www.ibm.com/support/ docview.wss?rs=3508&context=SS9NW2&uid=swg21385896.

   a. Create a folder in the **FastBackBMR** folder named **PEDrivers**.

   b. Create a folder in the **PEDrivers** folder named **StorageDrivers**.

   c. Copy all the downloaded drivers to the **StorageDrivers** folder.

The Windows PE software includes a predefined set of drivers. For any advanced customization procedures, ensure that you download the Windows 2008 drivers, so that they are compatible with the Microsoft Preinstallation Environment (PE).

5. (Optional) Download approved network drivers for Microsoft Windows PE. If you have systems with network cards that are not supported by the Windows 2008 PE, you have to add your own network drivers. For Microsoft PE to connect to the repository over the network, customize the CD image with network drivers that are suitable for your environment.

To download the approved network drivers, see http://www.ibm.com/ support/docview.wss?rs=3508&context=SS9NW2&uid=swg21385896. Download network drivers to match both your existing and future systems.

Copy the drivers to the `\FastBackBMR\PEDrivers\NetworkDrivers` folder by completing the following steps:

   a. Create a folder in the **FastBackBMR** folder named **PEDrivers**.

   b. Create a folder in the **PEDrivers** folder called **NetworkDrivers**

   c. Copy all the downloaded drivers to the **NetworkDrivers** folder.

6. (Optional) Download and install the data recovery utility, `TestDisk`, by completing the following steps.

   a. Download the latest version of `TestDisk` from this Web site: http://www.cgsecurity.org/wiki/TestDisk_Download.

   b. Save the download of the `TestDisk` image to your system. Complete the following steps:

      1) Run the Extraction Wizard by right-clicking the downloaded image and selecting **Extract All**.

      2) In the **Select a folder** page, type the work area directory, for example, `C:\FastBackBMR`.

      3) After the extraction is complete, rename the extracted folder (it is named something similar to `testdisk-6.9`) to TestDisk.

7. (Optional) Download and install the file viewer, `Lister`, by completing the following steps:

   a. Download the latest 32-bit version of `Lister` from this Web site: http://www.ghisler.com/lister/

   b. After the download is complete, extract the file to the designated storage area under `\FastBackBMR\Lister\`

Additional utilities, like **TestDisk** and **Lister**, can be downloaded and placed on the FastBack for Bare Machine Recovery CD. You can follow steps similar to steps 6 on page 10 and 7 on page 10. No add-on utilities are supported by IBM. For issues with add-on utilities, contact the vendor.

8. Create a FastBack for Bare Machine Recovery CD ISO image. Complete the following steps:

   a. Insert the *Tivoli Storage Manager for Bare Machine Recovery base code package* CD into your CD-ROM drive.

   b. Click **Start** > **Run**. Type cmd to open a command line window.

   c. Change to the CD-ROM drive where the FastBack for Bare Machine Recovery CD is located.

   d. Browse to the *Installation* folder.

   e. Run the following script:

      BuildTSMFastBackBMRCD.bat

   f. Type the path to the Microsoft AIK installation folder, for example, C:\Program Files\Windows AIK.

   g. Type the path to the work area, for example, C:\FastBackBMR.

   h. Type the path for the installation folder. The default path follows: C:\Program Files\Tivoli\TSM\FastBack BMR\bmr\FastBackWinBMR.

   i. Select the desired language to localize the FastBack for Bare Machine Recovery CD. Table 1 lists the languages that are supported by Tivoli Storage Manager FastBack for Bare Machine Recovery.

Table 1. Additional language support

Language	Language Code
Arabic (Saudi Arabia)	AR-SA
Czech	CS-CZ
Danish	DA-DK
German	DE-DE
Greek	EL-GR
English (United States)	EN-US
Spanish (Spain)	ES-ES
Finnish	FI-FI
French (France)	FR-FR
Hebrew	HE-IL
Hungarian	HU-HU
Italian	IT-IT
Japanese	JA-JP
Korean	KO-KR
Norwegian	NB-NO
Dutch	NL-NL
Polish	PL-PL
Portuguese (Brazil)	PT-BR
Portuguese (Portugal)	PT-PT
Russian	RU-RU
Swedish	SV-SE

Table 1. Additional language support (continued)

Language	Language Code
Turkish	TR-TR
Chinese (China)	ZH-CN
Chinese (Hong Kong)	ZH-HK
Chinese (Taiwan)	ZH-TW

j. If you select a language other than English (EN-US), you need to provide the path to the Windows AIK CD to copy the required language files. Type the drive letter of the Windows AIK CD.

   Note: If you only have one CD/DVD drive, copy the contents of the Windows AIK CD to a local drive, for example, `C:`.

k. Enter the desired PE keyboard layout.

l. The script creates an ISO image of the FastBack for Bare Machine Recovery CD under the output directory in the work area folder, for example, `C:\FastBackBMR\Output\winpe_x86.iso`.

m. Keep a copy of the log file that is reported at the completion of the build. You can use this log file to provide future verification of drivers included on the CD.

n. Burn this image to a CD to be used for a system recovery.

Chapter 3. Use Tivoli Storage Manager FastBack for Bare Machine Recovery

This section describes the tasks needed to use Tivoli Storage Manager FastBack for Bare Machine Recovery. Either boot your operating system from a FastBack for Bare Machine Recovery CD or use a non-system disk restore from FastBack Manager.

Using the FastBack for Bare Machine Recovery CD for bare machine recovery on your Windows system

If you want to boot your operating system from the FastBack for Bare Machine Recovery CD, follow the steps in this section. Follow these steps if you see a blue screen when you boot.

About this task

Specifically, if you see a blue screen, you can use the FastBack for Bare Machine Recovery CD to navigate to the `Any-to-Any HW Restore` utility. Using the `Any-to-Any HW Restore` utility, you can complete the following repair tasks:

- Create a backup of the current Windows status.
- Update storage drivers. For incorrect storage drivers, a blue screen with the following message is displayed: `STOP: 0x0000007B` . An incorrect jumper setting on an IDE or SCSI driver can also cause this stop error.
- Change the keyboard and mouse drivers.
- Change Hardware Abstraction Layer (HAL).
- Disable services.
- Disable AutoReboot.
- Remove upper and lower filters.

To use the FastBack for Bare Machine Recovery CD for bare machine recovery, complete the following steps:

Procedure

1. Insert the FastBack for Bare Machine Recovery CD into the drive of the failed system. This CD provides a bootable environment that you use to connect to the network, run Tivoli Storage Manager FastBack for Bare Machine Recovery, and connect to the backup storage location.
2. Select a network adapter. Most environments are able to receive DHCP configuration information so you do not have to manually type IP configuration information.
3. Click **Next**. The Remote FastBack repository Share Credentials window is displayed.
4. Type the Windows domain, user name, password, and Windows network share for the user who has access permissions to the backup storage location. The FastBack Server service does not need to run for the bare machine recovery to work.
5. Click **Next**.

6. Select a repository. Click **Refresh** to refresh the list of snapshots. If you click **Remove**, the selected repository is no longer displayed in the `Select repository` list. If you click **Close**, Tivoli Storage Manager FastBack for Bare Machine Recovery closes and a message is displayed. The message asks if you want to restart the system.

7. After you select a repository, select a system drive snapshot from the repository. Snapshots are defined by policy name, server, volume, and date.

8. Use one of the following procedures:
   - Option 1: Add a snapshot to restore one or more snapshots to a destination disk. By adding a snapshot, you create a partition table on the destination disk. The partition table is set up according to how many snapshots you are restoring. To add a snapshot, complete the following steps:

     a. Click **Add Snapshot**. The snapshot is displayed in the Bare Machine Recovery Snapshots section.

     b. Repeat the previous step for each partition on the target disk. You restore the additional snapshots to the same disk, but under a different partition on the target disk. The partition table is created based on the number of snapshots that you recover. For more information, see to the following table:

Table 2. Partition table details

Number of recovered snapshots	Partition table
1 snapshot	1 primary partition is created
2 snapshots	2 primary partitions are created
3 snapshots	3 primary partitions are created
4 snapshots	3 primary partitions and 1 extended partition with 1 logical drive are created
5 snapshots	3 primary partitions and 1 extended partition that has 2 logical drives are created

Each additional snapshot creates an additional logical drive within that last extended partition.

The order in which the snapshots are selected dictates the order in which they are restored to the destination disk. If the disk, that is being restored, had an EISA partition as the first partition, complete the following steps:

1) Select the volume with the EISA partition.

2) Select the volume with the operating system and boot information.

3) Select any additional volumes to restore.

Retain the original volume order to prevent unpredictable operating system behavior.

   c. Click **Restore**. The Bare Machine Recovery - Choose destination disk window is displayed.

   d. Select a destination disk that is large enough to meet the disk space requirements of the snapshots that are selected in steps 9a and 9b.

   e. For the **Bootable partition**, select the partition that you want to use. You can select **None**. If the first partition to be restored is an EISA partition, and you choose to make the first partition bootable, the selection is ignored and the next partition is made bootable.

Choose only one of the first three partitions as the primary boot partition. Other partitions are extended volumes, and Windows cannot boot from them.

Important: When you restore volumes from a Windows Server 2008 R2 operating system, you must restore the System Reserve volume in addition to restoring the operating volume. The System Reserve volume contains the boot loader that is required at boot time. The Windows Server 2008 R2 system fails to reboot if the System Reserve volume is not also restored. Choose the System Reserve as the bootable partition.

f. For **Disk signature**, select one of the following options:
   – **Use original disk signature**: The snapshot signature is used. This signature is the default value and it is uncommon to change this default.
   – **Use existing disk signature**: The selected destination disk signature is used.
   – **Set disk signature manually**: The **Selected signature** field is used to manually define the signature. The signature can contain up to eight hexadecimal characters.

g. View the values for required disk space. **Required disk space** indicates the amount of disk space that is required for the restore process. **Available disk space** is updated according to the selected disk. If the selected disk is too small, the field is highlighted with the color red.

h. Click **Start disk restore**. A summary is displayed. The summary includes the snapshots that are being restored. If there is not enough space on the destination disk, you cannot proceed.

i. Click **Yes** to proceed or **No** to go back to the Disk Selection window. If you click **Yes**, all data on the destination disk is erased.

j. Monitor progress for the restore process in the Bare Machine Recovery Snapshots section. Wait until all snapshots are marked as complete.

- Option 2: (Optional) Complete a volume restore to restore one, and only one, snapshot to the destination disk. If the drive is not bootable before the volume restore, the volume restore does not make the drive bootable. The restored snapshot is the operating system partition.

  This option is not intended for restoring to a new disk or new system. This volume restore option is to be used to restore the system volume of the original system with the original hardware. You are restoring the system volume backup on the original system volume.

  This option is to be used ONLY for restoring an existing system volume to a previous state that is using its snapshot. It is an uncommon scenario and a majority of users are NOT using this option.

  To compete a volume restore, complete the following steps:

  a. After you select a snapshot and policy, click the **Advanced...** button.
  b. Select the destination partition.
  c. Click **OK**.
  d. Restart the system.

  Complete the volume restore if you want to revert only the operating system volume to a previous version of the operating system. For example, if you upgraded to a new operating system version, but want roll back to the previous version of the operating system, use the volume restore. You

can also use the volume restore if an operating system volume is corrupted or infected by a virus, and you want to go to a previous snapshot of the operating system volume.

9. After the restore is complete, if the target system requires a different mass storage driver, specifically a mass storage driver that is different from the driver on the recovered snapshot, use the `Any-to-Any HW Restore` utility. For information about how to use the `Any-to-Any HW Restore` utility, see "Using the `Any-to-Any Hardware Restore` utility."

10. Start the restored system. If the system starts successfully, the Tivoli Storage Manager FastBack for Bare Machine Recovery process is complete.

If the system fails to start and you get a blue screen, the system hardware does not match the data structure for the snapshot. You can use the `Any-to-Any HW Restore` utility to recover data. For information about how to use the `Any-to-Any HW Restore` utility, see "Using the `Any-to-Any Hardware Restore` utility."

Using the `Any-to-Any Hardware Restore` utility
About this task

If you are restoring a system to hardware that is different from the original system, click **Any-to-Any HW Restore** and use one of the following procedures:

- If a black screen is displayed when you restart the system after the restore is complete, click **Repair Storage** and complete the following steps. **Repair Storage** is available only if the `TestDisk` tool is available.

  1. Select a storage device to repair.
  2. Click **Repair Storage Automatically**.
  3. Click **Exit Repair Storage**.

  You can repair more than one disk by using the Repair Storage tool. In case the automatic storage repair fails, click **Launch Test Disk** and repair the storage manually with the `TestDisk` tool. For information about how to use `TestDisk`, see http://www.cgsecurity.org/wiki/TestDisk.

- If a black screen is not displayed when you restart the system, click **Move to Step 2** and complete the following steps:

  1. In the Step 2 section, select the location of the Windows installation that requires repair.
  2. Verify that the correct Windows version is selected.
  3. In the Step 3 section, click **Create backup** to create a backup of the current Windows settings before the repair begins. This backup can be restored later by clicking **Restore Backup**.
  4. (Optional) In the Step 4 section, use one of the following procedures:
     - To select a specific driver to install, select either **Activate IDE drivers** or **Activate SCSI drivers**. This installs all the built-in IDE or SCSI drivers that exist in that version of Microsoft Windows (the version that is on the restored snapshot).
     - To install the generic driver, first download the required driver package to a USB flash drive. Then complete the following steps..
       a. Select **Generic Driver (Browse for a storage Driver)** from the driver list.
       b. Click **Install Driver**. A dialog box displays a list of available drivers.
       c. Insert the USB flash drive that contains the required driver package, then browse to and select the driver.

     d. Click **Open** to install the driver on the operating system you just restored with the bare machine recovery process.

     e. (Optional) You can disable and delete all previously installed SCSI/IDE drivers by selecting **Remove other SCSI/IDE Drivers**.

     f. After installing the drivers, you can click **Finish** or complete the optional next step.

5. (Optional) In the Step 5 section, set the following settings:

Change keyboard mouse drivers
> To reinstall keyboard and mouse drivers, select **Change keyboard mouse drivers** and click **Perform optional settings** to apply your selection.

Disable troublesome services
> To disable a predefined list of services, for example, UPS services, serial and parallel ports, select **Disable troublesome services**. You can also click **Services** to manually select services to disable.

Disable Upper/Lower filters
> Upper and lower filters are used by disk monitor applications, and can prevent Windows from running. In such cases, select **Disable Upper/Lower filters** to disable these filters and to run Windows. Click **Perform optional settings** to apply your selection.

Disable AutoReboot
> When a message like STOP: error is displayed, also known as a blue screen, and AutoReboot is enabled, the system is automatically restarted. To prevent this automatic restart, select **Disable AutoReboot** and click **Perform optional settings** to apply your selection.

Activate SYSPREP mini-setup
> Select this option to run the Windows mini-setup after the restart. Do not select this option on systems that serve as Active Directory Domain Controllers. The SYSPREP option removes the Active Directory from the system.

Services
> Click **Services** to manually select services that you want to disable.

HAL Settings
> Use caution when changing HAL settings. In some scenarios, you must reinstall the Hardware Abstraction Layer (HAL) to complete an any-to-any bare machine recovery process.
>
> When you click **HAL Settings**, information about the installed HAL is displayed. A list of alternative HALs is also displayed. To replace the current HAL, select a HAL from the list and click **Install HAL**. Some HAL conversions are not supported by Windows.

Perform optional settings
> Click **Perform optional settings** to edit the sysprep.inf file.

6. (Optional) Click **Restore Backup** to restore the backup you created in Step 3.

7. (Optional) Click **Exit program** to exit the Any-to-Any HW Restore utility without saving any information.

8. (Optional) Click **Finish** to save and close the Any-to-Any HW Restore utility.

Restoring a disk by using the non-system disk restore

This type of restore uses a system where FastBack Client is installed. To clone and deploy disks, the system is to have an extra disk or a USB disk.

About this task

When you restore a non-system disk from FastBack Manager, the FastBack for Bare Machine Recovery CD is not used for the restore.

To restore a disk or disks to different systems by using the non-system disk restore, complete the following steps:

Attention: If FastBack Client is not connected to the FastBack Server, this procedure cannot be completed.

Procedure

1. Before proceeding with a disk restore, close all open handles to the disk. For example, if it is open, close the Windows disk management utility on the target computer. This utility is part of Windows computer management tools.
2. Access FastBack Manager. Verify that the logon identity has privileges to run IBM Tivoli Storage Manager FastBack for Bare Machine Recovery for the relevant volumes.
3. Click **Monitor** and select the volumes to be restored by completing the following steps:

   Important: Retain the original disk layout to prevent unpredictable operating system behavior.
   a. (Optional) In the Monitor Filter section, select the volumes to be displayed and click **Filter**. In the display, click the **Start Time** column header to show the latest snapshots of the selected volumes.
   b. Right-click the snapshots to be restored. Use the shift key to select several snapshots. If the disk being restored has an EISA partition, do these steps:
      1) Select the volume with the EISA partition.
      2) Select the volume with the operating system and boot information.
      3) Select any additional volumes to restore.
4. Click **Bare Machine Recovery**.
5. You can complete the following tasks:
   - Change the order of the volumes by selecting a volume and clicking **Up** or **Down**.
   - If you want the restored disk to be bootable, check **Make first snapshot bootable**. Ensure that the first line contains the bootable volume. Do this taskto retain the original layout and order of the disk.

     Important: If the first partition to be restored is an EISA partition, and you choose to make the first partition bootable, the next partition is made bootable instead.
   - Check **Preserve disk signature** to preserve the disk structure and volume letters. This is important in boot disks.
   - Click **Select Destination Disk** to identify the location. There are two views available:

- **Tree View** - shows the servers and their disks, and in addition provides a folder (Standalone Disks) containing all the standalone disks
- **List View** - provides a summary view of all disks identified by the FastBack Server

- **Ignore open handles on the destination volume**

  If there are open handles on the target restore volume, the restore does not complete successfully. Open handles are any open files or applications that are running on the target restore volume. You can force a restore on a target volume with open handles by selecting **Ignore open handles on the destination volume** in the Restore Destination window. However, forcing a restore on a target volume with open handles can cause a problem with applications and loss of data in files that are open on the target volume.

- If you want to view properties for the disk, click **Disk properties**.

6. Click **Apply**. All data on the destination disk is destroyed. If one of the volumes is mounted and used by the application, the operation fails.

7. Click **OK** to restore the disk.

Scenarios using Tivoli Storage Manager FastBack

The following scenarios are common usage scenarios for Tivoli Storage Manager FastBack for Bare Machine Recovery.

About this task

These scenarios are listed in no particular order.

Scenario 1 - Restore complete bootable disk
> You need to restore a complete bootable disk. To complete this scenario, use the instructions in "Using the FastBack for Bare Machine Recovery CD for bare machine recovery on your Windows system" on page 13. In the instructions, the step that describes when you specify a bootable partition states that you can select *None*. Do not select *None*. After you select a bootable partition, you can click **Start Disk Restore**.

Scenario 2 - Dissimilar restore
> You need to restore more than one volume, including a boot volume, to a destination disk that is the boot disk. To complete this scenario, use the instructions in "Using the FastBack for Bare Machine Recovery CD for bare machine recovery on your Windows system" on page 13. Changes in the Active Directory might also be needed.

> If, after you complete the restore process, you restart the system and a blue screen is displayed, use the `Any-to-Any HW Restore` utility. For more information about the `Any-to-Any HW Restore` utility, see "Using the `Any-to-Any Hardware Restore` utility" on page 16.

Scenario 3 - Disk restore to a second disk that is a non-system disk
> You need to restore more than one volume, including a boot volume, to a destination disk that is not the boot disk. To complete this scenario, use the instructions in "Restoring a disk by using the non-system disk restore" on page 18. When we create and deploy a clone disk, remember that each operating system installation has unique system environment parameters.

> Change the following parameters manually to avoid a conflict with another system on the network. Check the following information before connecting the new system to the network, changing one, or more of the following parameters:

- Machine name
- Machine GUID
- IP address
- Subnet mask

Scenario 4 - Restore only system volume

You need to restore the boot volume to the same disk and the same volume. To complete this scenario, use the instructions in "Using the FastBack for Bare Machine Recovery CD for bare machine recovery on your Windows system" on page 13.

Scenario 5 - Restore from Tivoli Storage Manager server storage

You need to restore a FastBack DR Hub Server volume from Tivoli Storage Manager server storage. To complete this scenario, use the instructions in "Restoring from a Tivoli Storage Manager FastBack repository on a Tivoli Storage Manager server."

Using a single system for client and server

You can complete a bare machine recovery to a system that is used as a server and as a client.

When you start a bare machine recovery to a system that is used as a server and as a client, a message is displayed. The message indicates that the server is not accessible for synchronization. This message can be ignored. Click **Yes** to continue.

If you are going to complete a bare machine recovery to a system that serves as both FastBack Client and FastBack Server, use the following guidelines:

- For repositories, the following requirements must be met:
  - Do not use disk repositories. Volume, folder, or network share repositories can be used.
  - Do not locate the repository on the same disk as the system drive.
  - Do not back up the repository location.
  - You can have multiple repository locations on a dynamic disk or on a basic disk.
- The FastBack Client that is on the FastBack Server is configured as *SAN Enabled*. If the client is not *SAN Enabled*, performance is slower, there are problems with volume restore, and the disks can be displayed twice in the FastBack Manager disk display.

Restoring from a Tivoli Storage Manager FastBack repository on a Tivoli Storage Manager server

Use Tivoli Storage Manager FastBack for Bare Machine Recovery to recover data that is stored on a Tivoli Storage Manager FastBack repository replicated to a Tivoli Storage Manager server by using the FastBack Disaster Recovery.

Procedure

1. On the FastBack DR Hub Server, use the FBDRDir.exe and FBDRCopy.exe utilities to copy the FastBackBMR folder. This folder is on the Tivoli Storage Manager server in the REP_<branch_name>\REPOSITORY\DISK_xyz\ folder.
2. Use the node name and credentials associated with the Tivoli Storage Manager branch to copy the FastBackBMR folder.

3. Run `FBDRDir.exe` to get the branch names in the Tivoli Storage Manager server. This command lists all the REP_<branch_name> folders in the Tivoli Storage Manager server: `FBDRDir.exe -m; -f "TSM;TCPS:10.10.10.10;TCPP:1500;NODE:MyNodeName;PASSWORD:XXX;FILEPATH:*" -d 1`

4. Run `FBDRDir.exe` again and specify the required branch name and depth level 2: `FBDRDir.exe -m; -f "TSM;TCPS:10.10.10.10;TCPP:1500;NODE:MyNodeName;PASSWORD:XXX;FILEPATH:/rep_MyBranch" -d 2`

5. Use `FBDRCopy.exe` to copy the FastBackBMR folder to a local path: `FBDRCopy.exe -m; -f "TSM;TCPS:10.10.10.10;TCPP:1500;NODE:MyNodeName;PASSWORD:XXX;FILEPATH:/REP_MyBranch/REPOSITORY/DISK_E@0X09D6EC80/FastBackBMR" c:\tsm_bm`

6. Create a network share to the parent folder of the `FastBackBMR` folder.

7. Load the Tivoli Storage Manager FastBack for Bare Machine Recovery preinstallation environment on the system to be restored.

Using Tivoli Storage Manager FastBack for Bare Machine Recovery on your Linux system

Restore a Linux machine that is protected by the IBM Tivoli Storage Manager for FastBack client. A customized Linux image is used to boot the client machine to be recovered.

This procedure assumes the following criteria:

- The Linux boot CD recognizes the disk controllers, network adapters, and exposes the disk using Internet Small Computer System Interface (iSCSI) technology.
- A snapshot of a Linux root partition is available and it hosts a complete Linux system. The machine from which the snapshots were taken contains one disk with the following partition layout:

  `/boot /(root) swap`

- The Bare Machine Recovery machine contains the same hardware configuration as the machine that was backed up.

Note:
Tivoli Storage Manager FastBack for Bare Machine Recovery for Linux does not support a configuration that contains dissimilar hardware.

Creating a FastBack for Bare Machine Recovery on SUSE Linux

Part I: Use SUSE Studio to create a Linux Bootable CD to use with the FastBack for Bare Machine Recovery procedure:

1. Log in to your SUSE Studio account and click **Create new appliance**.
   a. In the `Choose a base template` section, select **Server** in the `Suse Linux Enterprise 11 SP1` field.
   b. In the `Select your architecture` section, select **32-bit**.
   c. Specify a name in the `Name your appliance` field and click **Create appliance**.

2. Click the **Software** tab.

3. Enter the word `scsi` in the `Search for Software` field. Scroll through the results and then click **+add** on the packages `lsscsi` and `iscsitarget`. Click the **Configuration** tab.

4. Select the `Default Locale` and `Default Timezone` according to your location. In the `Network` section, select `(DHCP)` unless you have a specific need for a static address. Clear **Enable Firewall** in the `Firewall` section. Click the **Build** tab.

5. Select **LiveCD** in the `Format` section and click the **Build** tab. Wait for the build process to complete.

6. Click **TestDrive** if you want to test the CD image before downloading it. Then click **Download**. If everything went fine, now you have an ISO image that you can use for performing BMR for a Linux box.

Part II: FastBack for Bare Machine Recovery using the SUSE Linux Bootable CD created in Part I:

1. Boot the machine that you want to restore from the SUSE Linux Bootable CD created in Part I.

2. Verify that you are using a valid IP address by running the `ipconfig` command. If a valid IP address is not available or a DHCP is not being used, configure the NIC by using Yast.

3. Launch Yast with the `yast` command.

4. In the Control Center, select the **Network Devices** <Tab> **Network Settings** <Enter>, and then proceed to the NIC configuration.

5. Display your IP address and make note of it:

```
# ifconfig

eth0 Link encap:Ethernet HWaddr 00:1e:8c:ff:31:51
inet addr:9.20.166.19 Bcast:9.20.166.255 Mask:255.255.255.0
inet6 addr: fe80::21e:8cff:feff:3151/64 Scope:Link [...]
```

6. List the disks available on the system by using the following command:

```
# lsscsi

[0:0:0:0] disk ATA Hitachi HDT72502 V5DO /dev/sda
[1:0:0:0] cd/dvd HL-DT-ST DVD-RAM GSA-H60N CX04 /dev/scd0
```

7. Expose the disk by using iSCSI target software by entering the following commands:

   Note: In case the disk that you want to restore contains an LVM partition, the iSCSI fails to expose the disk. You can check this using the following command:

```
linux:~ # fdisk -l /dev/sda
[]
Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *         1          13      104391   83  Linux
/dev/sda2            14        3916    31350847+  8e  Linux LVM
```

   If the disk you want to restore contains an LVM partition and you want to restore on that disk, run the following command:

```
linux:~ # vgchange -an
0 logical volume(s) in volume group "VolGroup00" now active
```

   Modify the file `/etc/ietd.conf` so that it contains the following lines at a minimum:

```
Target linuxbmr
Lun 0 Path=/dev/sda,Type=blockio
```

Where /dev/sda is the device name of the disk that you want to restore.

8. Restart the iSCSI target service with this command:

```
/etc/init.d/iscsitarget restart
```

9. Complete the tasks documented in "On the FastBack Server" on page 24.

Creating a FastBack for Bare Machine Recovery on Fedora

This procedure describes how to configure Fedora Live CD to use it when performing Bare Machine Recovery for Linux. This procedure assumes the following criteria:

- The ISO image of the Fedora Live CD was downloaded.
- The BMR Box has internet access. This is needed to download additional software that is not immediately available on the Live CD.

1. Create a CD of the ISO image and use it to boot the BMR Box.
2. Click Log In.
3. Enter the following commands:

```
su -
 yum install scsi-target-utils lsscsi
```

4. Enter y when prompted to download and install the following packages:

```
- lsscsi-0.23-1.fc12.i686.rpm
- perl-Config-General-2.44-1.fc12.noarch.rpm
- scsi-target-utils-1.0.1-0.fc12.i686.rpm
```

Create a FastBack for Bare Machine Recovery on Fedora with a live CD:

1. Start the iscsi target daemon with this command:

```
# service tgtd start
```

2. Create a new target with this command:

```
# tgtadm --lld iscsi --op new --mode target --tid=1 --targetname linuxbmr
```

3. Add the disk to be restored (/dev/sda in this procedure) to the new target:

```
# tgtadm --lld iscsi --op new --mode logicalunit --tid 1 --lun 1 -b /dev/sda
```

4. Make the target accessible through the network by issuing these two commands:

```
# tgtadm --lld iscsi --op bind --mode target --tid 1 -I ALL
```

```
# iptables -I INPUT -p tcp --dport 3260 -j ACCEPT
```

5. Complete the tasks documented in "On the FastBack Server" on page 24.

On the FastBack Server

1. Install the Microsoft iSCSI Software Initiator (if necessary).
2. Install Tivoli Storage Manager FastBack for Bare Machine Recovery.
3. Run the iSCSI initiator (**Start** → **All Programs** → **Microsoft iSCSI Initiator** → **Microsoft iSCSI Initiator**).
4. In the **Discover** tab, add a Target Portal by using the IP address previously assigned to the Linux Box.
5. In the **Targets** tab, select a target disk and log in to it.
6. Run FastBack Mount BMR for Linux using the new shortcut (**Start** → **All Programs** → **Tivoli Storage Manager** → **FastBack** → **FastBack BMR for Linux**)
7. Select the snapshot that contains the /boot partition, and then click **Add snapshot**.
8. Select the snapshot that contains the /partition, and then click **Add snapshot**.
9. Click **Restore**. Select the destination disk that you want to restore. The iSCSI disk is the last one added to the list. Ensure that a local disk is not selected as follows:
   a. Open the Computer Management window by selecting **Settings** → **Control Panel** → **Administrative Tools** → **Computer Management**.
   b. In the **Storage** branch, select **Disk Management**.
   c. Right-click each disk on the list and select **Properties**. For the iSCSI disk, the IET Virtual-Disk SCSI Device Disk Properties window is displayed. Ensure that you make a note of the disk number.
   d. Return to the **Targets** tab in FastBack Mount BMR for Linux and select the disk with the number noted in the previous step. Click **Start Disk Restore** and wait for the restore operation to complete.
   e. Verify that the disk state is online. If it is offline, right-mouse click the disk and select `Online`.
10. Close Tivoli Storage Manager FastBack for Bare Machine Recovery.
11. From the Microsoft iSCSI Initiator **Targets** tab, click **Details**, enable the check box in the Identifier area, and then select **Log off**. Verify that the status for the target is now inactive.
12. Complete the tasks documented in "On the Bare Machine Recovery machine."

On the Bare Machine Recovery machine

1. Stop the iSCSI target service. This will force all data to be written on the disk.
   SUSE Linux:

   ```
   # /etc/init.d/iscsitarget stop
   ```

   Fedora:

   ```
   # tgtadm --lld iscsi --op delete --mode logicalunit --tid 1 --lun 1
   # tgtadm --lld iscsi --op delete --mode target --tid=1
   # service tgtd stop
   ```

2. Run partprobe to tell the kernel to read the partition table again:

   ```
   # partprobe
   ```

Note: An `Invalid partition table` error might display when partprobe is issued for a volume that is not the volume that is being recovered. This is considered normal.

3. Make sure that the partition was restored:

```
# fdisk -l /dev/sda
Disk /dev/sda: 32.2 GB, 32212254720 bytes
255 heads, 63 sectors/track, 3916 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00018613

Device Boot  Start  End   Blocks    Id   System
/dev/sda1    1     13     104391    83   Linux
/dev/sda2    14    3661   29302560  83    Linux
```

4. Restore Grub bootloader so that you can boot the system:

```
# grub

grub> find /boot/grub/stage1
(or find /grub/stage1 if you get Error 15: File not found)
(hd0,0) #Grub is already installed on the first partition of first disk

grub> root (hd0,0) #Tell grub where / (root partition) can be found

grub> setup (hd0) #Tell grub where to install the bootload (this command will overwrite
                  the MBR)

grub> quit
```

- SUSE Linux or Fedora:

```
# mount /dev/sda1 /boot
# cd /boot
# grub
grub> find /grub/stage1 (hd0, 0)
#grub is already installed on the first partition of the disk
```

- On Fedora 17 you can use following commands to launch grub:

```
#mount /dev/sda2 /mnt
#/mnt/sbin/grub
```

If either /boot or /partition are restored on a different disk, the file `menu.lst` has to be updated.

- Fedora: The `menu.lst` file is in the /mnt/grub/ path.
- SUSE Linux: The `menu.lst` file is in the mnt/boot/grub/ path.

To edit it, the partition that contains `menu.lst` must be mounted:

```
# mount /dev/sda1 /mnt
```

Edit `menu.lst` for all the relevant sections:

```
root (hd0,0)
kernel /boot/vmlinuz-2.6.16.60-0.42.5-default root=/dev/sda2 vga=0x314
load_ramdisk=1 init=linuxrc resume=/dev/sda2 splash=silent showopts
initrd /boot/initrd-2.6.16.60-0.42.5-default
```

Unmount the partition:

```
# umount /mnt
```

5. Re-create a swap partition:

```
# fdisk /dev/sda

Command (m for help): n #Tell fdisk to create a new partition
Command action
l logical (5 or over)
p primary partition (1-4)
p
Partition Number (1-4): 3
First cylinder (3662-3916, default 3662): Enter
Using default value 3662
Last cylinder, +cylinders or +size{K,M,G} (3662-3916, default 3916): +2048M

Command (m for help): t # change partition type
Partition number (1-6): 3 # select partition we want to change
Hex code (type L to list codes): 82 # 82 stand for swap
Changed system type of partition 2 to 82 (Linux swap / Solaris)

Command (m for help): p # print partition table
[...]

Device  Boot   Start   End   Blocks    Id   System
/dev/sda1    1    13    104391    83    Linux
/dev/sda2    14   3661   29302560  83    Linux
/dev/sda3    3662  3916   2048287+  82    Linux swap / Solaris

Partition table entries are not in disk order

Command (m for help): w # save the changes
The partition table has been altered!
Calling ioctl() to re-read partition table.
WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table.
The new table will be used at the next reboot.
Syncing disks.
```

Initialize the swap space with this command:

```
# mkswap /dev/sda3
```

6. Mount the restored root volume:

```
# mount /dev/sda2 /mnt
```

7. Edit the swap configuration line in /mnt/etc/fstab to configure the partition that was created:

```
# vi /mnt/etc/fstab

/dev/sda3 none swap sw 0 0
```

8. Reboot the Linux box and remove the CD from the drive:

```
# reboot
```

Note: If you have a more complex partition layout, you might need to complete the tasks documented in "Restoring a more complex partition layout" on page 27.

Restoring a more complex partition layout

1. One disk with three partitions:

```
/boot
/
swap space
```

2. One disk with an LVM partition:

```
/boot
LVM partition hosting/
```

3. Two disks:

```
Disk 1: /boot and /
Disk 2: /opt (partition containing FastBack Client)
```

4. Restore more than one box at the same time from a hub server.

One disk with three partitions

In Step 7 of "On the FastBack Server" on page 24, select the snapshots to restore (both the /boot and / snapshots).

One disk with an LVM partition

This task is similar to the **One disk with three partitions** task. The logical volume is not re-created and the / partition is restored directly on the disk. The following example is an assumption about the configuration for backup and restore:

Original configuration:

```
Disk /dev/sda
Partition 1 /dev/sda1 /boot
Partition 2 /dev/sda2 VolGroup00
  VolGroup00
   LogVol00 /
   LogVol01 swap
```

Restored configuration:

```
Disk /dev/sda
Partition 1 /dev/sda1 /boot
Partition 2 /dev/sda2 /
Partition 3 /dev/sda3 swap
```

In Fedora and SUSE:

```
# fdisk /dev/sda

Command (m for help): p

Disk /dev/sda: 42.9 GB, 42949672960 bytes
255 heads, 63 sectors/track, 5221 cylinders, total 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000dde27

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *          63      208844      104391   83  Linux
/dev/sda2          208845    79638476    39714816   8e  Linux LVM
/dev/sda3        79638477    83886079     2123801+  82  Linux swap / Solaris

Command (m for help): t
Partition number (1-4): 2
Hex code (type L to list codes): 83
Changed system type of partition 2 to 83 (Linux)

Command (m for help): p

Disk /dev/sda: 42.9 GB, 42949672960 bytes
255 heads, 63 sectors/track, 5221 cylinders, total 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000dde27

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *          63      208844      104391   83  Linux
/dev/sda2          208845    79638476    39714816   83  Linux
/dev/sda3        79638477    83886079     2123801+  82  Linux swap / Solaris

Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table.
WARNING: Re-reading the partition table failed with error 16: Device or resource
busy. The kernel still uses the old table.
The new table will be used at the next reboot.
Syncing disks.
```

Before rebooting from the Live CD as described in Step 9 of "On the Bare
Machine Recovery machine" on page 24, these extra steps are required:

1. Edit the boot parameters:

```
# mount /dev/sda1 /mnt
# vi /mnt/grub/menu.lst
```

2. Edit the kernel line from

```
kernel /vmlinuz-2.6.18-164.6.1.el5 ro
root=/dev/VolGroup00/LogVol00
```

to

```
kernel /vmlinuz-2.6.18-164.6.1.el5
ro root=/dev/sda2
 save and close
# umount /mnt
```

so that it points to the new root partition.

3. Edit fstab to point to the new disks:

```
# mount /dev/sda2 /mnt
# vi /mnt/etc/fstab
```

4. Change from

```
/dev/VolGroup00/LogVol00    /                        ext3    defaults  1 1
LABEL=/boot              /boot                    ext3    defaults     1 2
Tmpfs     /dev/shm      tmpfs    defaults        0 0
 Devpts    /dev/pts              devpts  gid=5,mode=620  0 0
sysfs                    /sys                      sysfs    defaults  0 0
proc                     /proc                    proc     defaults      0 0
/dev/VolGroup00/LogVol01 swap                      swap     defaults      0 0
```

to

```
/dev/sda2       /                    ext3    defaults       1 1
LABEL=/boot                  /boot                    ext3    defaults     1 2
tmpfs                        /dev/shm                 tmpfs    defaults    0 0
devpts                       /dev/pts                 devpts  gid=5,mode=620  0 0
sysfs                        /sys                     sysfs    defaults      0 0
proc                         /proc                    proc     defaults      0 0
/dev/sda3     swap                    swap     defaults       0 0

save and close

# umount /mnt
```

5. Restart BMR:

```
# reboot
```

Two disks

In Step 7 of "Creating a FastBack for Bare Machine Recovery on SUSE Linux" on page 21, change the iSCSI configuration in the following way:

```
Target LinuxBMR
Lun 0 Path=/dev/sda,Type=blockio
Lun 1 Path=/dev/sdb,Type=blockio
```

In Steps 9 and 10 of "On the FastBack Server" on page 24, repeat these steps for each disk that must be restored.

Restore more than one box at a time from a hub server

This example assumes that there are *n* machines to restore. On each machine, modify these steps:

1. In Step 7 of "Creating a FastBack for Bare Machine Recovery on SUSE Linux" on page 21, change the iSCSI configuration in the following way:

```
Target LinuxBMR_n
Lun 0 Path=/dev/sda,Type=blockio
```

2. In Steps 9 and 10 of "On the FastBack Server" on page 24, repeat these steps for each disk that must be restored.

Troubleshooting tips:

1. If after the last reboot, the operating system does not boot and the grub> prompt displays, repeat Step 4 to configure the grub boot loader again.

2. If your machine has an IDE disk and the original `fdisk -l` command displayed /dev/hda, the bootable CD (Fedora or SUSE), might display `fdisk -l` as /dev/sda. This is an incorrect device name (`sda` instead of `hda`). This incorrect device name can affect situations where the /etc/fstab or grub.menu files must be edited. Make sure to use the original disk name (/dev/hda).

3. If the iSCSI fails to log off after a bare machine recovery for Linux, you can manually stop the iSCSI.

```
# /etc/init.d/tgtd status
tgtd (pid 2271 2269) is running...
# kill -9 2271
# /etc/init.d/tgtd status
tgtd (pid 2269) is running...
# kill -9 2269
# /etc/init.d/tgtd status
tgtd dead but subsys locked
# /etc/init.d/tgtd stop
# /etc/init.d/tgtd start
Starting SCSI target daemon: [ OK ]
# /etc/init.d/tgtd stop
Stopping SCSI target daemon: [ OK ]
```

Chapter 4. Troubleshooting

This information describes some common problems that you might have with the Tivoli Storage Manager FastBack for Bare Machine Recovery and provides possible solutions.

Problem

When you boot from the FastBack for Bare Machine Recovery CD by using a static IP address, the IP address is not resolved. This problem occurs when the system used to resolve the IP address is a Windows Internet Name Service server (WINS). In the IP configuration window, you configure the client IP address and DNS server address, but there are no fields to help you configure the WINS.

Solution

Use IP addresses when you connect to the repository, or manually configure the WINS. To manually configure the WINS, complete the following steps:

1. When the IP address configuration window is displayed, click **Next**.

2. In the window that requests credentials, click the icon for information.

3. Click **Command prompt**.

4. In the command prompt window, add the WINS. The following string is an example of how to add the WINS:

```
netsh interface ipv4 add winsserver "local area
 connection" [WINS IP address]
```

Problem

When downloading the Windows Automated Installation Kit (AIK) from Microsoft, you are not sure what version to download.

Solution

There are two versions of the Windows Automated Installation Kit (AIK):

- The old version labeled Windows Automated Installation Kit (AIK). This version was posted in 2007.

- A current version labeled Automated Installation Kit (AIK) for Windows Vista SP1 and Windows Server 2008. This version was posted in 2008.

Use the current version of the Windows AIK.

Problem

After you download the Windows Automated Installation Kit (AIK), you are not sure what to do next.

Solution

If the downloaded file is an IMG file, change the file extension to ISO. Burn the ISO file content to a CD or mount with an ISO mounting utility. You can then install the Windows Automated Installation Kit (AIK).

Problem

How are mass storage drivers loaded?

Solution

Select "Browse for a storage driver" and click "Install Driver". In the dialog box that opens, browse for the driver you want to install. After you complete the driver installation, click "Finish" and close the application to restart the system.

Problem

During the Tivoli Storage Manager FastBack for Bare Machine Recovery build process, an error occurs.

Solution

Verify that there is sufficient space to build the ISO image. In addition, verify that the image is built on a Windows 32-bit operating system.

Problem

What are the requirements for booting in a VMware environment?

Solution

If you are booting in a VMware environment, make sure that you allocate enough memory to meet the memory requirements on the guest VMware system. If you have less memory on the system, the user interface for the software does not display.

For more information about supported environments, including memory requirements, see "x86 processors (32-bit) with Windows 32-bit operating systems" on page 2.

Problem

When using the FastBack for Bare Machine Recovery CD for bare machine recovery, a network adapter cannot be selected.

Solution

The network adapter cannot be identified. See step 5 on page 10 to verify that the correct network drivers needed for the particular network adapter are downloaded and placed in the `\FastbackBMR\PEDrivers\NetworkDrivers` folder.

Problem

After completing information in the remote FastBack repository share credentials window, click **Next**. An error message is displayed. When using the FastBack for Bare Machine Recovery CD for bare machine recovery and selecting the network adapter, the IP address was not updated to match the network adapter.

Solution

After you select the network adapter, update the IP address. If this update does not occur, or, if the Windows domain, user name, password, and path to the repository share are not configured correctly, the system cannot access the FastBack Server repository and a message is displayed. If the message code is FBSP7024, the FastBack for Bare Machine Recovery CD was created on a Windows 64-bit system. You need to rebuild the FastBack for Bare Machine Recovery CD on a Windows 32-bit operating system.

If the message code is FBSP7008E, open a command prompt window and complete the following steps:

1. Enter the `netsh` command to manually configure the network adapter IP configuration. For example:

   ```
   netsh interface ipv4 set address "Local Area Connection" static
    192.168.100.145 255.255.255.0 192.168.100.1
   ```

2. Enter the `ipconfig` command to confirm the Windows IP configuration settings. For example:

   ```
   ipconfig
   ```

3. Enter the `net use` command to map a network drive to the repository share. For example, if the repository share *FB_REP_E* is on a system at the IP address *192.168.100.100*, issue the following command:

```
net use M: \\192.168.100.100\FB_REP_E
```

A message stating the command was completed successfully is displayed. Confirm that the network drive is available by entering the following command:

```
net use
```

4. Manually start the `Fastbackmount.exe` program in bare machine recovery mode by entering the following command:

```
M:\FastBackBMR\mount\Fastbackmount -bmr -load M: user=Administrator
 domain=TSMFBDOMAIN password=xxxxx -pe 5.5.3.0 -log X:\
```

M: represents the drive letter for the repository. Continue with step 6 on page 14 and identify the network share name that you created when selecting the repository share in step 3 on page 32.

If FastBack Mount does not start from the command line, you can start FastBack Mount from the Windows Task Manager by completing the following steps:

a. Enter the `net use` command. Include the repository drive and IP address in the command. For example:

```
net use M: \\192.168.100.1\FB_REP_E
```

A message reports that the password is not valid. Enter the user name and password.

b. Enter the `taskmgr` command. For example:

```
taskmgr
```

c. In the Windows Task Manager window, click **New Task**.

d. Enter the following command:

```
M:\FastBackBMR\mount\Fastbackmount -bmr -load M: user=Administrator
 domain=TSMFBDOMAIN password=xxxxx -pe 5.5.3.0 -log X:\
```

If errors occur, additional information about the error is displayed. For example, if there are mismatched Windows system DLL files, a message is displayed.

Note: Errors are reported to the command prompt and logged to the `.SF` file that is stored at the root of the `X:\RAM-DISK` drive.

Problem

When using the FastBack for Bare Machine Recovery CD for bare machine recovery and choosing a destination disk, you do not see any disk devices listed.

Solution

The necessary storage driver was not loaded because of a missing or incorrect Windows driver. Go to the `FastBackBMR/StorageDrivers` directory and look at the contents of the `PEDrivers` folder. You created this folder when you created the FastBack for Bare Machine Recovery CD. Confirm that the necessary driver is in the folder. For more information about downloading approved storage drivers for Microsoft Windows PE, see step 4 on page 9.

Appendix. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in the Tivoli Storage Manager family of products:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

The Tivoli Storage Manager Information Center and related publications are enabled for accessibility. For information about the accessibility features of the information center, see the following topic: http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4/topic/com.ibm.help.ic.doc/iehs36_accessibility.html. Read about the accessibility features of the information center.

Keyboard navigation

On Windows, the Tivoli Storage Manager product family follows Microsoft conventions for all keyboard navigation and access. Drag-and-drop support is managed by using the Microsoft Windows accessibility option known as *MouseKeys*. For more information about MouseKeys and other Windows accessibility options, see the Windows online help, citing the keyword "MouseKeys".

On other operating systems, these products follow the operating-system conventions for keyboard navigation and access.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (http://www.ibm.com/able) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Adobe is either a registered trademark or a trademark of Adobe Systems Incorporated in the United States, and/or other countries.

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

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

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

Other product and service names might be trademarks of IBM or other companies.

Index

Glossary

This glossary includes terms and definitions.

To view glossaries for other IBM products, go to http://www.ibm.com/software/globalization/terminology.

The following cross-references are used in this glossary:
- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode
> In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)
> In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode
> An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment
> The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See *access control list*.

activate
> To validate the contents of a policy set and then make it the active policy set.

active-data pool
> A named set of storage pool volumes that contain only active versions of client backup data.

active file system
> A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set
> The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version
The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log
A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup
A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client
A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule
A database record that describes the planned processing of an administrative command during a specific time period. See also *client schedule*.

administrative privilege class
See *privilege class*.

administrative session
A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

administrator
A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

Advanced Program-to-Program Communication (APPC)
An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

agent node
A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate
An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

aggregate data transfer rate
A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

APPC See *Advanced Program-to-Program Communication*.

application client
A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

archive

> To copy programs, data, or files to other storage media, usually for long-term storage or security. Contrast with *retrieve*.

archive copy

> A file or group of files that was archived to server storage.

archive copy group

> A policy object containing attributes that control the generation, destination, and expiration of archived files.

archive-retention grace period

> The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

association

> (1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.
>
> (2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

audit To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication

> The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

authentication rule

> A specification that another user can use to either restore or retrieve files from storage.

authority

> The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

> A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

> A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

> See *automounted file system*.

automatic detection

> A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B

backup-archive client

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bind   To associate all versions of a file with a management class name. See *rebind*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

C

cache   To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD   See *client acceptor*.

central scheduler
A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor
An HTTP service that serves the applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX®, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon* (CAD).

client acceptor daemon (CAD)
See *client acceptor*.

client domain
The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node
A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session
A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file
An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set
A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode
A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule
A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server
Pertaining to the model of interaction in distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file
A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the dsm.sys file. See also *client user-options file*.

client user-options file
A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the dsm.opt file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration
A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation
The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group
A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point
A point in time when data is considered consistent.

Common Programming Interface for Communications (CPI-C)
A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication method
The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

communication protocol
A set of defined interfaces that permit computers to communicate with each other.

compression
A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager
A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation
A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup
A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted

copy group
A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool
A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D

daemon
A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file
A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode
A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series
One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot
A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication
A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server
A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover
A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)
A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class
A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication
See *data deduplication.*

demand migration
The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client
The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination
A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class
A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file
(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver
A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)
A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan
A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI

See *data storage-management application-programming interface*.

dynamic serialization

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

enterprise logging

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process

prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list
See *include-exclude list*.

execution mode
A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration
The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file
A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend
To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)
Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

extent The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

external library
A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSLS). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F

file access time
On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age
For migration prioritization purposes, the number of days since a file was last accessed.

file device type
A device type that specifies the use of sequential access files on disk storage as volumes.

file server
A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space
A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore,

retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)
> A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state
> The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)
> A kernel extension that intercepts all file system operations and provides any space management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state
> The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency
> A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup
> The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup
> A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy
> A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G

General Parallel File System
> A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)
> In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state
> The state of all file systems to which space management has been added when space management is globally deactivated for a client node. When

space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system.

GPFS™
See *General Parallel File System*.

GPFS node set
A mounted, defined group of GPFS file systems.

group backup
The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)
A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client
A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client
See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup
A backup of a full file system or raw logical volume as a single object.

inactive file system
A file system for which space management has been deactivated. Contrast with *active file system*.

inactive version
A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file
A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list
A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup
(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore
See *mailbox restore*.

information lifecycle management (ILM)
GPFS policy-based file management for storage pools and file sets.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the IP standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer
See *LAN-free data movement*.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.

(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local (1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.

(2) For HSM products, pertaining to the destination of migrated files that are being moved.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy
The space that is used by logical files in a storage pool. This space does

not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)
An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)
In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume
A portion of a physical volume that contains a file system.

logical volume backup
A backup of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)
Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)
A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU See *logical unit*.

LUN See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

macro file
A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore
A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object
In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, server definitions, and server group definitions.

managed server
A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class
A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit
The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB   See *megabyte*.

media server
In a z/OS® environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)
(1) 1 048 576 bytes (2 to the 20th power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk storage capacity and communications volume, 1 000 000 bits.

metadata
Data that describes the characteristics of data; descriptive data.

migrate
To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file
A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode
A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job
A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold
High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring
The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode
In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit
The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point
On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period
The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period
The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N

Nagle algorithm
An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe
A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node
A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system
A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format
A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP

    See *Network Data Management Protocol*.

NetBIOS

    See *Network Basic Input/Output System*.

network-attached storage (NAS) file server

    A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)

    A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)

    A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

    A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node   A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name

    A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

    A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

non-native data format

    A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode

    A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O

offline volume backup

    A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

    A backup in which the volume is available to other system applications during the backup operation.

open registration

    A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class

    A privilege class that gives an administrator the authority to disable or halt

the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file
A file that contains processing options. On Windows and NetWare systems, the file is called dsm.opt. On AIX, UNIX, Linux, and Mac OS X systems, the file is called dsm.sys.

originating file system
The file system from which a file was migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file
A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode
A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing
In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode
A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation
A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (`passwordaccess` option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character
See *wildcard character*.

physical file
A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy
The amount of space that is used by physical files in a storage pool. This

space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in
A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain
A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class
A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set
A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file
A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database
A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named `.SpaceMan` in each file system to which space management has been added.

premigration
The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage
A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool
A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class
A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also *bind*.

recall In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode

A mode that is assigned to a migrated file with the **dsmattr** command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold
   The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation
   The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log
   A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register
   To define a client node or administrator ID that can access the server.

registry
   A repository that contains access and configuration information for users, systems, and software.

remote
   (1) Pertaining to a system, program, or device that is accessed through a communication line.

   (2) For HSM products, pertaining to the origin of migrated files that are being moved.

resident file
   On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore
   To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention
   The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve
   To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back
   To remove changes that were made to database files since the last commit point.

root user
   A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN   See *storage area network*.

schedule
A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode
The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume
A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)
A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup
The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration
The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall
The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

serialization
The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

server A software program or a computer that provides services to other software programs or other computers.

server options file
A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode
A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

server storage
The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session
> A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

session resource usage
> The amount of wait time, processor time, and space that is used or retrieved during a client session.

shared dynamic serialization
> A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

shared library
> A library device that is used by multiple storage manager servers.

shared static serialization
> A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

snapshot
> An image backup type that consists of a point-in-time view of a volume.

space-managed file
> A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

space management
> The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

space manager client
> A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

space monitor daemon
> A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file
> A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file
> On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL   See *Secure Sockets Layer*.

stabilized file space
A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window
A time period during which a schedule must be initiated.

static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent
A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)
A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy
(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool
A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume
A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class
A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows

transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that can be used to eliminate the need to recall a migrated file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

Systems Network Architecture (SNA)

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T

tape library

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

threshold migration

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode

A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can include substitution for command parameters and conditional logic.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name
A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)
The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V

validate
To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version
A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space
A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume
An archive file on a target server that represents a sequential media volume to a source server.

volume
A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file
A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service
A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup
A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore
A function that uses a Microsoft Volume Shadow Copy Service (VSS)

software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore
A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup
A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore
A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

W

wildcard character
A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation
A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name
A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)
A partition within a single operating system instance.

IBM®

Product Number: 5724-U95

Printed in USA