IBM System Storage SAN Volume Controller

Host Attachment User's Guide Errata

Version 4.1.x

SC26-7905-00-Errata

Note:
Before using this information and the product it supports, read the information in **Notices**.

Contents

About this guide

This guide provides errata information that pertains to version 4 release 1 of the IBM System Storage SAN Volume Controller Host Attachment User's Guide.

This guide contains the corrections and additions of the SAN Volume Controller Host Attachment User's Guide on a per chapter basis.

Who should use this guide

This guide is intended for system administrators or others who install and use the SAN Volume Controller.

Before using the SAN Volume Controller, you should have an understanding of storage area networks (SANs), the storage requirements of your enterprise, and the capabilities of your storage units.

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

Boldface	Text in **boldface** represents menu items and command names.
Italics	Text in *italics* is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values, such as a default directory or the name of a cluster.
`Monospace`	Text in `monospace` identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

Last update

This guide was last updated November 17, 2006.

Chapter 1. Open-systems hosts

You can attach the SAN Volume Controller to open-systems hosts that use the small computer system interface-fibre channel protocol (SCSI-FCP). You can also attach the SAN Volume Controller to iSCSI (small computer system interface over internet protocol) hosts that are attached to a Cisco fabric.

The SAN Volume Controller supports connection to the Cisco MDS 9000 SAN-OS Software Release 2.1 for the Cisco MDS 9000 Family platform with the attached iSCSI hosts (in single path mode only) shown in Table 1.

Table 1. Overview of the Cisco MDS Family platform

Operating System	Drivers	Notes
Windows 2000 (32 bit only) (SP4)	Microsoft initiator driver rev 1.05a and Microsoft initiator driver rev 1.06	Clustering is not supported.
Windows 2003 (32 bit only) (SP1)	Microsoft initiator driver rev 1.06	
Red Hat EL Advanced Server 2.1	Linux initiator rev 3.6.2	
Red Hat EL Advanced Server 3.0		

Note: Guidelines for zoning iSCSI hosts and SAN Volume Controller:

In a conventional fibre-channel SAN, there are a number of SAN paths between a particular SAN Volume Controller I/O group and the server HBA ports that use the VDisks that are supplied by that I/O group. A multipathing device driver resolves these multiple paths into a single logical device on which the server can perform I/O. The multipathing device driver also provides failover and path recovery functions for SAN fabric paths that change or fail.

The present iSCSI solution, however, only supports a single path between the iSCSI host network information center (NIC) and the SAN Volume Controller VDisk. There is no multipathing driver in the iSCSI host. Therefore, there is no recovery from errors and it is not possible to concurrently upgrade the SAN Volume Controller firmware while maintaining connectivity with an iSCSI host system. It is, thus, inappropriate for the SAN Volume Controller to present the VDisk to multiple ports in the fibre-channel SAN. To prevent this, you must select a single SAN Volume Controller port in each SAN Volume Controller I/O group that is to be associated with each iSCSI host. Zoning is then applied in the Cisco MDS switch so that each iSCSI host can only see one SAN Volume Controller port in each SAN Volume Controller I/O group. If multiple iSCSI hosts are in use, the hosts should be evenly spread across the ports in each SAN Volume Controller I/O group. Issue the `svctask mkvdiskhostmap` command to ensure that each SAN Volume Controller VDisk is mapped to a single NIC in the server.

Hosts are attached to the SAN Volume Controller through a switched fibre-channel fabric. Each SAN Volume Controller fibre-channel adapter has two ports, and each port is identified by a worldwide port name (WWPN).

SAN Volume Controller does not limit the number of fibre-channel ports or host bus adapters (HBAs) that each connected host or host partition can have. Your connected hosts are limited only by the number of ports or HBAs that are supported by the multipathing device driver on the host (or host partition).

Table 2 provides an overview of the host systems that are supported by the SAN Volume Controller.

Table 2. Overview of supported host systems

Brand	Server host systems	Operating system
HP	HP 9000	HP-UX 11i
		HP-UX 11.0
	HP AlphaServer	Tru64 UNIX®
IBM®	eServer™ BladeCenter® JS20	AIX®
		Linux®
	eServer i5	AIX
	eServer iSeries™	AIX
	eServer p5	AIX
	eServer pSeries®	AIX
		Linux
	eServer pSeries SP™	AIX
	eServer zSeries®	Linux
	RS/6000®	AIX
	RS/6000 SP	AIX
	System p5™	AIX
	System z9™	Linux
SGI	Origin 3000 Series	IRIX
	Origin 350	IRIX
Sun	SPARC servers	Solaris
various	Intel-based and AMD-based	Linux
		Microsoft® Windows® 2000 Server
		Microsoft Windows 2000 Advanced Server
		Microsoft Windows 2003 Server
		Microsoft Windows NT®
		Novell NetWare
VMware	VMware ESX	various guest operating systems

The following IBM Web site provides current interoperability information about current support information, including maximum configuration details, technical

flashes, hints and tips, host systems, operating system levels, HBAs, cables, fabrics that IBM supports, and documentation about the SAN Volume Controller:

http://www.ibm.com/storage/support/2145

Chapter 2. Attaching to IBM pSeries and JS20 hosts running the Linux operating system

This information provides an overview for attaching the SAN Volume Controller to supported POWER™ technology-based hosts running the Linux operating system.

The following POWER technology-based hosts are supported by the SAN Volume Controller:

- IBM eServer pSeries
- IBM eServer BladeCenter JS20
- IBM System p5

Setting queue depth with QLogic HBAs

Configure your host running the Linux operating system using the formula specified in Chapter 1 of the SAN Volume Controller Configuration Guide. The queue depth is the number of I/O operations that can be run in parallel on a device.

Perform the following steps to set the maximum queue depth:

1. Add the following line to the /etc/modules.conf file:

   For the 2.4 kernel (SUSE Linux Enterprise Server 8 or Red Hat Enterprise Linux 3):

   ```
   options qla2300 ql2xmaxqdepth=new_queue_depth
   ```

   For the 2.6 kernel (SUSE Linux Enterprise Server 9 or Red Hat Enterprise Linux 4):

   ```
   options qla2xxx ql2xmaxqdepth=new_queue_depth
   ```

2. Rebuild the RAM disk that is associated with the kernel being used by using one of the following commands:
   - If you are running on an SUSE Linux Enterprise Server operating system, run the `mk_initrd` command.
   - If you are running on a Red Hat operating system, run the `mkinitrd` command and then restart.

Setting queue depth for Emulex HBAs

Configure your host running the Linux operating system to allow a maximum queue depth of four.

Perform the following steps to set the maximum queue depth:

1. Add the following line to the /etc/modules.conf file:

   ```
   options lpfc lpfc_lun_queue_depth=new_queue_depth
   ```

   where `new_queue_depth` is a number calculated using the formula specified in Chapter 1 of the SAN Volume Controller Configuration Guide.
2. Restart the machine.

Chapter 3. Attaching to IBM System z9 or eServer zSeries hosts running the Linux operating system

This information provides an overview for attaching the SAN Volume Controller to supported IBM System z9™ and IBM eServer zSeries hosts running the Linux operating system.

Known issues and limitations for System z9 and zSeries hosts

There are some restrictions for System z9 and zSeries hosts running the Linux operating system.

The following Web site provides currently known restrictions for the latest Linux for System z9 and zSeries streams:

http://www.ibm.com/developerworks/linux/linux390/

From that Web site, click **June 2003 stream** to find the known restrictions for the Linux kernel 2.4, or click **October 2005 stream** to find known restrictions for the Linux kernel 2.6.

Interoperability restrictions with Red Hat Enterprise Linux 4 update 4 for IBM System z

There are some restrictions specific to SAN Volume Controller V4.1.0 interoperation with Red Hat Enterprise Linux 4 Update 4 for IBM System z™.

Additional restrictions might be imposed on hardware, such as switches and storage, that are attached to SAN Volume Controller.

Installation restrictions

In order to install Red Hat Enterprise Linux 4 Update 4 onto a SAN Volume Controller FCP device, you must have at least one DASD device connected to the system via ESCON® or FICON®; otherwise, the installation will fail.

IPL restrictions

DM-MP multipathing is not available on either the root or boot devices.

For more information about DM-MP multipath usage, see http://www.redhat.com/docs/manuals/csgfs/browse/rh-cs-en/ap-rhcs-dm-multipath-usagetxt.html. System re-IPL (`shutdown -r`) is supported on zVM guests only; not in LPAR mode.

Multipath configuration

Red Hat Enterprise Linux 4 Update 4 does not include a default multipath configuration for SAN Volume Controller.

You must update the device part of your multipath.conf with the following:

```
device {
               vendor                  "IBM "
               product                 "2145           "
               path_grouping_policy    group_by_prio
```

```
        prio_callout            "/sbin/mpath_prio_alua /dev/%n"
        features                "1 queue_if_no_path"
        path_checker            tur
}
```

Fabric maintenance

You need to apply a workaround on the host before you can begin fabric maintenance.

Apply the following workaround on the Red Hat Enterprise Linux 4 Update 4 zlinux host before undertaking fabric maintenance, including the SAN Volume Controller code load:

```
$>vi /bin/bug23366.sh
          #!/bin/bash
          for f in `grep offline /sys/bus/scsi/drivers/sd/*/state | sed 's/^\(.*state\):.*$/\1/g'`;
          do echo running > $f;
          done
$>crontab -e
          */2 * * * * /bin/bug23366
```

Chapter 4. Attaching to a NetApp V-Series or gFiler NAS server

This information provides an overview for attaching the SAN Volume Controller to a NetApp V-Series or gFiler NAS server (NetApp).

Attachment requirements for NetApp servers

This section provides an overview of the requirements for attaching the SAN Volume Controller to a NetApp V-Series or gFiler server.

The requirements for attaching the SAN Volume Controller to your NetApp server running the Network Appliance Data ONTAP operating system:

- Check the LUN limitations for your NetApp server. Ensure that there are enough fibre-channel adapters installed in the server to handle the total LUNs that you want to attach.
- Ensure that you have the documentation for your NetApp server and the *IBM System Storage SAN Volume Controller: Installation Guide*. All SAN Volume Controller publications are available from the following Web site:

  http://www.ibm.com/storage/support/2145
- Ensure that you have installed the correct operating system level on your NetApp server.

Installing the HBA and driver on a NetApp server

The NetApp servers are supplied with preinstalled host bus adapters (HBAs). If additional HBAs are required, contact your NetApp service representative for advice on which model of HBA to install.

The Data ONTAP installation on your NetApp server includes the HBA driver, so no special installation steps are necessary for the HBA driver.

Configuring the Data ONTAP software on a NetApp server

You must configure the Data ONTAP software before you can use NetApp servers with the SAN Volume Controller.

The NetApp server requires an external root volume, which you can create with either of the following methods:

- Create a VDisk on your SAN Volume Controller and map it to your NetApp server.
- Partition and zone a back-end storage controller so that your NetApp server can directly access a suitable volume to use as its root volume.

Before you configure the Data ONTAP software, an IBM service representative must have installed the SAN Volume Controller.

After the prerequisite task is complete, use the following general steps to configure your Data ONTAP software:

1. Zone the NetApp server to the SAN Volume Controller on the fibre-channel SAN. Ensure that exactly two paths exist between the NetApp server and each I/O group on the SAN Volume Controller. For redundancy, configure the switch zoning so that host bus adapter (HBA) port A in the NetApp server is zoned with a single connection to SAN Volume Controller node A in an I/O Group, while HBA port B in the NetApp server is zoned with a single connection to SAN Volume Controller node B in the same I/O group. When you use a SAN Volume Controller cluster with multiple I/O Groups, each HBA port in the NetApp server should be zoned to one SAN Volume Controller node in each I/O Group.
2. Create the host system on the SAN Volume Controller, using the worldwide port names (WWPNs) of the HBAs in the NetApp server. Map the VDisks to the host as required.
3. Create aggregates and volumes on your NetApp server using the instructions in your host system publications.

Managing VDisks with NetApp servers

Before you manage your virtual disks (VDisks) on your NetApp server, you must consider some important issues.

The following information is important when managing your VDisks with a NetApp server:

- If you use the *–fmtdisk* option (or the corresponding GUI option) to create a formatted VDisk on a SAN Volume Controller that is to be mapped to a NetApp server, wait until the format operation completes before creating the host mapping to associate that VDisk with the NetApp server.
- The NetApp server does not support shrinkage or expansion of VDisks. Shrinkage is not possible, but to achieve the same effect as expansion, you can perform the following steps:
  1. Create a new VDisk on the SAN Volume Controller.
  2. Map the new VDisk to the NetApp server.
  3. Use the NetApp management tools to add the new VDisk to the desired NetApp aggregate.

Limitations and restrictions when using NetApp servers

Before you use your NetApp server, consider some limitations and restrictions.

Review the following limitations and restrictions when using NetApp servers:

1. You cannot use SAN Volume Controller Copy Services (FlashCopy, Metro Mirror, and Global Mirror) to copy VDisks that are mapped to NetApp servers. This limitation applies only to VDisks that are mapped to NetApp and does not restrict the use of Copy Services on other VDisks.
2. The maximum VDisk size supported by NetApp servers is 500 GB, which equates to 500x1024x1024x1000 bytes. However, the *minimum* VDisk size supported by NetApp servers is 1 GB, which equates to 1024x1024x1024 bytes. The definition for 1 GB used in SAN Volume Controller is 1024x1024x1024 bytes, so mapping a 1GB SAN Volume Controller VDisk to a NetApp server will work, but mapping a 500 GB SAN Volume Controller VDisk to a NetApp server will fail.

3. Vdisks that are mapped to NetApp servers can be moved between I/O Groups on SAN Volume Controller, but you must halt the NetApp server before you do this.
4. You must not map VDisks to a NetApp server as LUN 0. This is the default behavior when creating a host mapping on SAN Volume Controller, and you must override this by using the -scsi switch on the mkvdiskhostmap command.
5. You can import pre-existing NetApp LUNs into SAN Volume Controller in image mode, except in the case of the NetApp server's root volume. If SAN Volume Controller is introduced into an existing NetApp installation, either:
   - the NetApp root file system must be rebuilt using a new VDisk presented by SAN Volume Controller.
   - the NetApp server's root file system must remain on the original controller, directly accessed by the NetApp server (and masked from SAN Volume Controller by, for example, LUN Partitioning or switch zoning).
6. The NetApp server and SAN Volume Controller might share a back-end storage controller if:
   - appropriate LUN Partitioning is in place on the back-end storage controller
   - the back-end controller does *not* implement "LUN Grouping" (where all LUNs in a group must be accessed through the same storage controller, and changing the controller used to access one LUN causes the other LUNs in the group to move to that same controller)
   - the back-end controller is supported by both the NetApp server and by SAN Volume Controller

Chapter 5. Attaching to an SGI Origin host running the SGI IRIX operating system

This information provides the requirements and other information for attaching the SAN Volume Controller to a Silicon Graphics (SGI) Origin host running the SGI IRIX operating system.

Attachment requirements for SGI Origin hosts

This section provides an overview of the requirements for attaching the SAN Volume Controller to an SGI Origin server running the IRIX operating system.

The requirements for attaching the SAN Volume Controller to your SGI Origin host system running the IRIX operating system are as follows:

- Check the LUN limitations for your host system. Ensure that there are enough fibre-channel adapters installed in the server to handle the total LUNs that you want to attach.
- Ensure that you have the documentation for your host system and the *IBM System Storage SAN Volume Controller: Installation Guide*. All SAN Volume Controller publications are available from the following Web site:

  http://www.ibm.com/storage/support/2145
- Ensure that you have installed the correct operating system level and any updates.
- Review device driver installation documents and configuration utility documents for additional patches that you might need.

Environments for SGI Origin hosts

Ensure that your SGI Origin host uses a supported operating system and version.

The SAN Volume Controller supports SGI Origin hosts that run the IRIX Operating System 6.5.16 Release and above.

The following IBM Web site provides current interoperability information about supported software levels:

http://www.ibm.com/storage/support/2145

HBAs for SGI Origin hosts

Ensure that your SGI Origin hosts use the correct host bus adapters (HBAs).

The SAN Volume Controller supports SGI Origin hosts running the IRIX operating system that use QLogic HBAs.

The following IBM Web site provides current interoperability information about supported HBAs:

http://www.ibm.com/storage/support/2145

Drivers and firmware for SGI Origin hosts

Be sure that you use the correct host bus adapter device driver and firmware levels for SGI Origin hosts running on the IRIX operating system.

The IRIX operating system includes the QLogic HBA driver, so no special installation steps are necessary for the QLogic HBA driver. The following IBM Web site provides current interoperability information about device driver and firmware levels:

http://www.ibm.com/storage/support/2145

Installing the HBA on an SGI Origin host

The first step for attaching the SGI Origin host is to install the host bus adapter (HBA).

Before you install the HBA, ensure that the adapter is supported by the SAN Volume Controller. See the supported hardware list at the following IBM Web site if you need to verify that the HBA is supported:

http://www.ibm.com/storage/support/2145

To install the HBA, use the following general steps:

1. Shut down your host and its attached peripherals, following the manufacturer's recommendations.
2. Install the HBA, using the manufacturer's installation instructions.

Configuring the QLogic HBA for SGI Origin hosts

After you install the QLogic host bus adapter (HBA) and driver, you must configure the HBA.

XVM Volume Manager failover capability

The SAN Volume Controller supports XVM failover capability for SGI Origin hosts.

The *XVM Volume Manager Administrator's Guide* describes the configuration and administration of XVM logical volumes. See Chapter 4 for information about XVM failover.

The SAN Volume Controller supports version 2 of XMV failover, so you need to create and edit the /etc/failover2.conf file.

To set up the SGI host, complete the following steps:

1. Rescan the HBA ports: `scsiha -rp <device>`.
2. Find the physical paths of volumes within XVM: `show -v *`.
3. Create SGI labels and partitions on the volume: `/usr/bin/fx -x -d <physical path>`.
4. Create /etc/failover2.conf manually. For HBA load balancing, choose different paths.
5. Either reboot the SGI host or initialize failover.
6. Label volumes within XVM: `label -name <labelname> <path>`.
7. Create slices and volumes within XVM: `slice -volname <volname> /phys/<name>`.

8. Create the xfs filesystem on the volumes: `mkfs -t xfs <path>`.
9. Create mount directories.
10. Mount the volumes.
11. Update /etc/fstab.

```
#lun0_svc
 /dev/dsk/5005076801000deb/lun0vol/c4p400000 affinity=0
 /dev/dsk/5005076801000deb/lun0vol/c3p200000 affinity=0
 /dev/dsk/5005076801000df8/lun0vol/c3p100000 affinity=1
 /dev/dsk/5005076801000df8/lun0vol/c4p300000 affinity=1
#lun1_svc
 /dev/dsk/5005076801000deb/lun1vol/c3p100000 affinity=0
 /dev/dsk/5005076801000deb/lun1vol/c4p300000 affinity=0
 /dev/dsk/5005076801000df8/lun1vol/c4p400000 affinity=1
 /dev/dsk/5005076801000df8/lun1vol/c3p200000 affinity=1
```

Figure 1. An example of failover2.conf

To display, configure, or change the settings for the XVM physical volumes:

- Use the XVM hardware inventory command to display the actual status for preferred / alternate paths: `hinv –c disk`
- Use the XVM foconfig command to parse the failover2.conf file on a running system and configure the settings for the preferred or alternate path.
- Use the XVM foswitch command to change the settings for the preferred or alternate path and access a physical volume.

SAN boot support on SGI Origin hosts

SGI does not support SAN boot for SGI Origin hosts that run the IRIX operating system.

Known issues and limitations

There are some known issues and limitations for attaching the SAN Volume Controller to an SGI Origin host.

The attached SGI Origin host running the IRIX operating system has the following limitations:

- SGI's failover infrastructure does not include multipath load balancing.
- SGI does not support SAN boot in this environment.
- The IRIX command scsifo -d to show Failover status does not work properly.
- The IRIX command scsifo -s to initiate manual failover does not work properly.

The following IBM support Web site provides for the most current information about known restrictions:

http://www.ibm.com/storage/support/2145

Chapter 6. Fibre-channel port name identification

This is an overview of the fibre-channel port name identification for the following host systems:

- HP 9000
- HP AlphaServer
- IBM System p5™, eServer, or RS/6000
- Linux
- NetApp
- SGI
- Sun
- Windows 2000 and Windows 2003
- Windows NT
- VMware

The WWPN consists of 16 hexadecimal characters (0 - 9 and A - F). The SAN Volume Controller uses it to uniquely identify the fibre-channel HBA that is installed in your host system. The SAN Volume Controller automatically finds the WWPN for your host fibre-channel HBA when you attach your host system to the SAN Volume Controller.

Note: If your host uses more than one fibre-channel HBA to connect to your SAN Volume Controller, you must add multiple entries to the host list for this host. You must add one for each fibre-channel HBA. Each HBA will have a unique WWPN.

The format and content of the fibre-channel port identifier are determined by the manufacturer of the link control facility for the applicable fibre-channel port. The identifier is an eight-byte field, which the fibre-channel protocols use to uniquely identify the fibre-channel port.

Locating the WWPN for a NetApp server

You can locate the WWPN for a NetApp server by following the steps in this topic.

1. Start the NetApp server.
2. At the NetApp system console, run the following command: `sysconfig -v`.

   Figure 2 on page 18 shows an example of the command output where the WWPNs are 500a098200004060 and 500a098300004060.

```
netapp_system_console> sysconfig -v
<snip>
        slot 0: FC Host Adapter 0a (Dual-channel, QLogic 2322 rev. 3, 64-bit, N-port,<UP>)
                Firmware rev:   3.3.220
                Host Port Id:   0x690913        FC Node Name:   5:00a:098200:004060
<snip>
        slot 0: FC Host Adapter 0b (Dual-channel, QLogic 2322 rev. 3, 64-bit, N-port,<UP>)
                Firmware rev:   3.3.220
                Host Port Id:   0x640913        FC Node Name:   5:00a:098300:004060
<snip>
netapp_system_console>
```

Figure 2. An example of the sysconfig command output

Locating the WWPN for an SGI Origin host

You can locate the WWPN for an SGI Origin host running the IRIX operating system with a QLogic adapter by following the steps in this topic.

1. Restart the server.
2. Type the `scsiha -w [bus_number | device]` command. For example, type `scsiha -w 6 7 8 9`. Figure 3 shows an example of the command output.

```
# scsiha -w 6 7 8 9
6 Portname: 210000e08b05d207
7 Portname: 210000e08b04d539
8 Portname: 210000e08b050808
9 Portname: 210000e08b038fe6
#
```

Figure 3. An example of the scsiha — bus_number device | command

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 give 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.

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 PUBLICATIONS "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.

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 measurement may have been

estimated through extrapolation. Actual results may vary. Users of this document may 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 may be addressed to the suppliers of those products.

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

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.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

- Enterprise Storage Server
- FlashCopy
- IBM
- IBM eServer
- IBM TotalStorage
- IBM System Storage
- System Storage
- TotalStorage

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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

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

Glossary

This glossary includes terms for the IBM System Storage SAN Volume Controller.

This glossary includes selected terms and definitions from A Dictionary of Storage Networking Terminology (http://www.snia.org/education/dictionary), copyrighted 2001 by the Storage Networking Industry Association, 2570 West El Camino Real, Suite 304, Mountain View, California 94040-1313. Definitions derived from this book have the symbol (S) after the definition.

The following cross-references are used in this glossary:

See Refers the reader to one of two kinds of related information:

- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
- A synonym or more preferred term.

See also
Refers the reader to one or more related terms.

Contrast with
Refers the reader to a term that has an opposite or substantively different meaning.

access mode
One of three different modes in which a logical unit (LU) in a disk controller system can operate. See also *image mode*, *managed space mode*, and *unconfigured mode*.

auxiliary virtual disk
The virtual disk that contains a backup copy of the data and that is used in disaster recovery scenarios. See also *master virtual disk*.

cluster
In SAN Volume Controller, a pair of nodes that provides a single configuration and service interface.

configuration node
A node that acts as the focal point for configuration commands and manages the data that describes the cluster configuration.

connected
In a Global Mirror relationship, pertaining to the status condition that occurs when two clusters can communicate.

Copy Services
The services that enable you to copy virtual disks (VDisks): FlashCopy, Metro, and Global Mirror.

counterpart SAN
A nonredundant portion of a redundant storage area network (SAN). A counterpart SAN provides all the connectivity of the redundant SAN but without the redundancy. Each counterpart SANs provides an alternate path for each SAN-attached device. See also *redundant SAN*.

cross-volume consistency
In SAN Volume Controller, a consistency group property that guarantees

consistency between virtual disks when an application issues dependent write operations that span multiple virtual disks.

data migration
The movement of data from one physical location to another without disrupting I/O operations.

degraded
Pertaining to a valid configuration that has suffered a failure but continues to be supported and legal. Typically, a repair action can be performed on a degraded configuration to restore it to a valid configuration.

dense wavelength division multiplexing (DWDM)
A technology that places many optical signals onto one single-mode fiber using slightly different optical frequencies. DWDM enables many data streams to be transferred in parallel.

dependent write operations
A set of write operations that must be applied in the correct order to maintain cross-volume consistency.

destage
A write command initiated by the cache to flush data to disk storage.

device

- In the CIM Agent, the storage server that processes and hosts client application requests.
- IBM definition: A piece of equipment that is used with the computer and does not generally interact directly with the system, but is controlled by a controller.
- HP definition: In its physical form, a magnetic disk that can be attached to a SCSI bus. The term is also used to indicate a physical device that has been made part of a controller configuration; that is, a physical device that is known to the controller. Units (virtual disks) can be created from devices after the devices have been made known to the controller.

directed maintenance procedures
The set of maintenance procedures that can be run for a cluster. These procedures are run from within the SAN Volume Controller application and are documented in the service guide.

disconnected
In a Metro or Global Mirror relationship, pertains to two clusters when they cannot communicate.

discovery
The automatic detection of a network topology change, for example, new and deleted nodes or links.

disk controller
A device that coordinates and controls the operation of one or more disk drives and synchronizes the operation of the drives with the operation of the system as a whole. Disk controllers provide the storage that the cluster detects as managed disks (MDisks).

disk drive
A disk-based, nonvolatile, storage medium.

disk zone
A zone defined in the storage area network (SAN) fabric in which the SAN Volume Controller can detect and address the logical units that the disk controllers present.

empty In a Global Mirror relationship, a status condition that exists when the consistency group contains no relationships.

ESS See *IBM TotalStorage® Enterprise Storage Server®*.

exclude
To remove a managed disk (MDisk) from a cluster because of certain error conditions.

excluded
In SAN Volume Controller, the status of a managed disk that the cluster has removed from use after repeated access errors.

extent A unit of data that manages the mapping of data between managed disks and virtual disks.

fabric port (F_port)
A port that is part of a fibre-channel fabric. An F_port on a fibre-channel fabric connects to the node port (N_port) on a node.

failover
In SAN Volume Controller, the function that occurs when one redundant part of the system takes over the workload of another part of the system that has failed.

FlashCopy® mapping
A relationship between two virtual disks.

FlashCopy relationship
See *FlashCopy mapping*.

FlashCopy service
In SAN Volume Controller, a copy service that duplicates the contents of a source virtual disk (VDisk) to a target VDisk. In the process, the original contents of the target VDisk are lost. See also *point-in-time copy.*

F_port See *fabric port*.

fuzzy copy
In Global Mirror, the copy that is provided when running in the asynchronous mode. The target virtual disk (VDisk) is not necessarily consistent with the source VDisk at every point in time. The host application writes data to the source VDisk and receives the final status on the write operation before the data is actually written to the target VDisk.

Global Mirror
An asynchronous copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk that is designated in the relationship.

grain In a FlashCopy bitmap, the unit of data represented by a single bit.

HBA See *host bus adapter*.

HLUN
See *virtual disk*.

host An open-systems computer that is connected to the SAN Volume Controller through a fibre-channel interface.

host bus adapter (HBA)
In SAN Volume Controller, an interface card that connects a host bus, such as a peripheral component interconnect (PCI) bus, to the storage area network.

host ID
In SAN Volume Controller, a numeric identifier assigned to a group of host fibre-channel ports for the purpose of logical unit number (LUN) mapping. For each host ID, there is a separate mapping of Small Computer System Interface (SCSI) IDs to virtual disks (VDisks).

host zone
A zone defined in the storage area network (SAN) fabric in which the hosts can address the SAN Volume Controllers.

hub A fibre-channel device that connects nodes into a logical loop by using a physical star topology. Hubs will automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

IBM TotalStorage Enterprise Storage Server (ESS)
An IBM product that provides an intelligent disk-storage subsystem across an enterprise.

ID See *identifier*.

identifier (ID)
A sequence of bits or characters that identifies a user, program device, or system to another user, program device, or system.

idle In a FlashCopy relationship, the state that occurs when the source and target virtual disks (VDisks) act as independent VDisks even if a mapping exists between the two. Read and write caching is enabled for both the source and the target.

idling

- The status of a pair of virtual disks (VDisks) that have a defined copy relationship for which no copy activity has yet been started.
- In a Metro or Global Mirror relationship, the state that indicates that the master virtual disks (VDisks) and auxiliary VDisks are operating in the primary role. Consequently, both VDisks are accessible for write I/O operations.

idling-disconnected
In a Global Mirror relationship, the state that occurs when the virtual disks (VDisks) in this half of the consistency group are all operating in the primary role and can accept read or write I/O operations.

illegal configuration
A configuration that will not operate and will generate an error code to indicate the cause of the problem.

image mode
An access mode that establishes a one-to-one mapping of extents in the managed disk (MDisk) with the extents in the virtual disk (VDisk). See also *managed space mode* and *unconfigured mode*.

image VDisk
A virtual disk (VDisk) in which there is a direct block-for-block translation from the managed disk (MDisk) to the VDisk.

inconsistent
In a Metro or Global Mirror relationship, pertaining to a secondary virtual disk (VDisk) that is being synchronized with the primary VDisk.

inconsistent-copying
In a Global Mirror relationship, the state that occurs when the primary virtual disk (VDisk) is accessible for read and write input/output (I/O) operations, but the secondary VDisk is not accessible for either. This state occurs after a **start** command is issued to a consistency group that is in the inconsistent-stopped state. This state also occurs when a **start** command is issued, with the force option, to a consistency group that is in the idling or consistent-stopped state.

inconsistent-disconnected
In a Global Mirror relationship, a state that occurs when the virtual disks (VDisks) in the half of the consistency group that is operating in the secondary role are not accessible for either read or write I/O operations.

inconsistent-stopped
In a Global Mirror relationship, the state that occurs when the primary virtual disk (VDisk) is accessible for read and write input/output (I/O) operations, but the secondary VDisk is not accessible for either read or write I/O operations.

input/output (I/O)
Pertaining to a functional unit or communication path involved in an input process, an output process, or both, concurrently or not, and to the data involved in such a process.

interoperability
The capability to communicate, run programs, or transfer data among various functional units in a way that requires the user to have little or no knowledge of the unique characteristics of those units.

I/O See *input/output.*

I/O group
A collection of virtual disks (VDisks) and node relationships that present a common interface to host systems.

I/O throttling rate
The maximum rate at which an I/O transaction is accepted for this virtual disk (VDisk).

managed disk (MDisk)
A Small Computer System Interface (SCSI) logical unit that a redundant array of independent disks (RAID) controller provides and a cluster manages. The MDisk is not visible to host systems on the storage area network (SAN).

managed disk group
A collection of managed disks (MDisks) that, as a unit, contain all the data for a specified set of virtual disks (VDisks).

managed space mode
An access mode that enables virtualization functions to be performed. See also *image mode* and *unconfigured mode*.

mapping
See *FlashCopy mapping*.

master virtual disk
The virtual disk (VDisk) that contains a production copy of the data and that an application accesses. See also *auxiliary virtual disk.*

MDisk
See *managed disk*.

mesh configuration
A network that contains a number of small SAN switches configured to create a larger switched network. With this configuration, four or more switches are connected together in a loop with some of the paths short circuiting the loop. An example of this configuration is to have four switches connected together in a loop with ISLs for one of the diagonals. The SAN Volume Controller does not support this configuration.

Metro Mirror
A synchronous copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk that is designated in the relationship.

migration
See *data migration*.

node One SAN Volume Controller. Each node provides virtualization, cache, and Copy Services to the storage area network (SAN).

node name
A name identifier associated with a node. (SNIA)

node port (N_port)
A port that connects a node to a fabric or to another node. N_ports connect to fabric ports (F_ports) or to other N_ports of other nodes. N_ports handle creation, detection, and flow of message units to and from the connected systems. N_ports are end points in point-to-point links.

node rescue
In SAN Volume Controller, the process by which a node that has no valid software installed on its hard disk drive can copy the software from another node connected to the same fibre-channel fabric.

N_port
See *node port*.

NWWN
See *worldwide node name*.

offline
Pertaining to the operation of a functional unit or device that is not under the continual control of the system or of a host.

online Pertaining to the operation of a functional unit or device that is under the continual control of the system or of a host.

operating set
In SAN Volume Controller, the set of nodes that are operating together to deliver storage services.

oversubscription
The ratio of the sum of the traffic that is on the initiator N-node connections to the traffic that is on the most heavily loaded Inter-Switch Links (ISLs), where more than one ISL is connected in parallel between these switches. This definition assumes a symmetrical network and a

specific workload that is applied equally from all initiators and sent equally to all targets. See also *symmetrical network*.

partition

- IBM definition: A logical division of storage on a fixed disk.
- HP definition: A logical division of a container represented to the host as a logical unit.

partner node
The other node that is in the I/O group to which this node belongs.

partnership
In Metro or Global Mirror, the relationship between two clusters. In a cluster partnership, one cluster is defined as the local cluster and the other cluster as the remote cluster.

PLUN See *managed disk*.

point-in-time copy
The instantaneous copy that the FlashCopy service makes of the source virtual disk (VDisk). In some contexts, this copy is known as a T_0 *copy*.

port The physical entity within a host, SAN Volume Controller, or disk controller system that performs the data communication (transmitting and receiving) over the fibre channel.

port ID
An identifier associated with a port.

prepared
In a Global Mirror relationship, the state that occurs when the mapping is ready to start. While in this state, the target virtual disk (VDisk) is offline.

preparing
In a Global Mirror relationship, the state that occurs when any changed write data for the source virtual disk (VDisk) is flushed from the cache. Any read or write data for the target VDisk is discarded from the cache.

primary virtual disk
In a Metro or Global Mirror relationship, the target of write operations issued by the host application.

PWWN
See *worldwide port name*.

queue depth
The number of I/O operations that can be run in parallel on a device.

quorum disk
A managed disk (MDisk) that contains quorum data and that a cluster uses to break a tie and achieve a quorum.

quorum index
The pointer that indicates the order used to resolve a tie. Nodes attempt to lock the first quorum disk (index 0), followed by the next disk (index 1), and finally the last disk (index 2). The tie is broken by the node that locks them first.

rack A free-standing framework that holds the devices and card enclosure.

redundant SAN
A storage area network (SAN) configuration in which any one single component might fail, but connectivity between the devices within the

SAN is maintained, possibly with degraded performance. This configuration is normally achieved by splitting the SAN into two, independent, counterpart SANs. See also *counterpart SAN.*

rejected
A status condition that describes a node that the cluster software has removed from the working set of nodes in the cluster.

relationship
In Metro or Global Mirror, the association between a master virtual disk (VDisk) and an auxiliary VDisk. These VDisks also have the attributes of a primary or secondary VDisk. See also *auxiliary virtual disk, master virtual disk, primary virtual disk, and secondary virtual disk.*

reliability
The ability of a system to continue to return data even if a component fails.

remote fabric
In Global Mirror, the storage area network (SAN) components (switches and cables) that connect the components (nodes, hosts, and switches) of the remote cluster.

roles Authorization is based on roles that map to the administrator and service roles in an installation. The switch translates these roles into SAN Volume Controller administrator and service user IDs when a connection is made to the node for the SAN Volume Controller.

SAN See *storage area network*.

SAN Volume Controller fibre-channel port fan in
The number of hosts that can see any one SAN Volume Controller port.

SATA See *Serial Advanced Technology Attachment*.

SCSI See *Small Computer Systems Interface*.

SCSI back-end layer
The layer in a Small Computer Systems Interface (SCSI) network that performs the following functions: controls access to individual disk controller systems that are managed by the cluster; receives requests from the virtualization layer, processes them, and sends them to managed disks; addresses SCSI-3 commands to the disk controller systems on the storage area network (SAN).

SCSI front-end layer
The layer in a Small Computer Systems Interface (SCSI) network that receives I/O commands sent from hosts and provides the SCSI-3 interface to hosts. SCSI logical unit numbers (LUNs) are mapped to virtual disks (VDisks) in this layer as well. Thus, the layer converts SCSI read and write commands that are addressed to LUNs into commands that are addressed to specific VDisks.

SDD See *subsystem device driver (SDD).*

secondary virtual disk
In Metro or Global Mirror, the virtual disk (VDisk) in a relationship that contains a copy of data written by the host application to the primary VDisk.

sequential VDisk
A virtual disk that uses extents from a single managed disk.

Serial Advanced Technology Attachment (SATA)
The evolution of the ATA interface from a parallel bus to serial connection architecture. (S)

SLP See *Service Location Protocol*.

Small Computer System Interface (SCSI)
A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

SMI-S See *Storage Management Initiative Specification*.

SNIA See *Storage Networking Industry Association*.

stand-alone relationship
In FlashCopy, Metro Mirror, and Global Mirror, relationships that do not belong to a consistency group and that have a null consistency group attribute.

stop A configuration command that is used to stop the activity for all copy relationships in a consistency group.

stopped
The status of a pair of virtual disks (VDisks) that have a copy relationship that the user has temporarily broken because of a problem.

storage area network (SAN)
A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. (S)

Storage Management Initiative Specification (SMI-S)
A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface that allows storage management systems to identify, classify, monitor, and control physical and logical resources in a storage area network. The interface is intended as a solution that integrates the various devices to be managed in a storage area network (SAN) and the tools used to manage them.

Storage Networking Industry Association (SNIA)
An association of producers and consumers of storage networking products whose goal is to further storage networking technology and applications. See www.snia.org.

striped
Pertains to a virtual disk (VDisk) that is created from multiple managed disks (MDisks) that are in the MDisk group. Extents are allocated on the MDisks in the order specified.

subsystem device driver (SDD)
An IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

superuser authority
The level of access required to add users.

suspended
The status of a pair of virtual disks (VDisks) that have a copy relationship that has been temporarily broken because of a problem.

switch A network infrastructure component to which multiple nodes attach. Unlike hubs, switches typically have internal bandwidth that is a multiple of link bandwidth, and the ability to rapidly switch node connections from one to another. A typical switch can accommodate several simultaneous full link bandwidth transmissions between different pairs of nodes. (S) Contrast with *hub*.

symmetrical network
A network in which all the initiators are connected at the same level and all the controllers are connected at the same level.

synchronized
In Metro or Global Mirror, the status condition that exists when both virtual disks (VDisks) of a pair that has a copy relationship contain the same data.

system
A functional unit, consisting of one or more computers and associated software, that uses common storage for all or part of a program and also for all or part of the data necessary for the execution of the program. A computer system can be a stand-alone unit, or it can consist of multiple connected units.

topology
The logical layout of the components of a computer system or network and their interconnections. Topology deals with questions of what components are directly connected to other components from the standpoint of being able to communicate. It does not deal with questions of physical location of components or interconnecting cables. (S)

trigger
To initiate or reinitiate copying between a pair of virtual disks (VDisks) that have a copy relationship.

unit identifiers (UIDs)
A unit identifier can be one of the following:

1. an integer expression whose value must be zero or positive
2. an * (asterisk) that corresponds to unit 5 for input or unit 6 for output
3. the name of a character array, character array element, or character substring for an internal file

unconfigured mode
A mode in which I/O operations cannot be performed. See also *image mode* and *managed space mode*.

unmanaged
An access mode that pertains to a managed disk (MDisk) that is not used by the cluster.

valid configuration
A configuration that is supported.

VDisk See *virtual disk*.

virtual disk (VDisk)
In SAN Volume Controller, a device that host systems attached to the storage area network (SAN) recognize as a Small Computer System Interface (SCSI) disk.

virtualized storage
Physical storage that has virtualization techniques applied to it by a virtualization engine.

virtual storage area network (VSAN)
A fabric within the SAN.

vital product data (VPD)
Information that uniquely defines system, hardware, software, and microcode elements of a processing system.

VLUN See *managed disk*.

VSAN See *virtual storage area network*.

worldwide node name (WWNN)
An identifier for an object that is globally unique. WWNNs are used by Fibre Channel and other standards.

worldwide port name (WWPN)
A unique 64-bit identifier associated with a fibre-channel adapter port. The WWPN is assigned in an implementation- and protocol-independent manner.

WWNN
See *worldwide node name*.

WWPN
See *worldwide port name*.

zoning
In fibre-channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones.

Index

S

T

V

W

X

Z

Readers' Comments — We'd Like to Hear from You

IBM System Storage SAN Volume Controller
Host Attachment User's Guide Errata
Version 4.1.x

Publication No. SC26-7905-00-Errata

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name	Address
Company or Organization	
Phone No.	E-mail address

Readers' Comments — We'd Like to Hear from You
SC26-7905-00-Errata

Cut or Fold
Along Line

Fold and Tape **Please do not staple** Fold and Tape

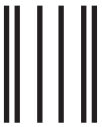

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9032 South Rita Road
Tucson, Arizona 85755-4401

Fold and Tape **Please do not staple** Fold and Tape

Cut or Fold
Along Line

Printed in USA