

IBM TotalStorage SAN Volume Controller



Planning Guide

Version 2.10

IBM TotalStorage SAN Volume Controller



Planning Guide

Version 2.10

Fifth Edition (February 2005)

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

© Copyright International Business Machines Corporation 2003, 2005. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	v
Tables	vii
About this guide	ix
Who should use this guide?	ix
Summary of changes	ix
Summary of changes for GA22-1052-04 SAN Volume Controller Planning Guide	ix
Summary of changes for GA22-1052-03 SAN Volume Controller Planning Guide	x
Emphasis	x
SAN Volume Controller library and related publications	xi
Related Web sites	xii
How to order IBM publications	xiii
How to send your comments	xiii
Chapter 1. Virtualization and the IBM TotalStorage SAN Volume Controller	1
Virtualization	1
The need for virtualization	3
Fabric level virtualization models	3
Symmetric virtualization	4
SAN Volume Controller	5
Uninterruptible power supply overview	8
Overview of backup functions	11
Cluster configuration backup	12
Chapter 2. Installation planning	15
Preparing your SAN Volume Controller environment	15
Power cables for the uninterruptible power supply 5115	16
Power cables for the uninterruptible power supply 5125	17
Preparing your uninterruptible power supply environment	19
Ports and connections	20
Chapter 3. Preparing the physical configuration	21
Completing the hardware location chart	21
Hardware location guidelines	22
Hardware location chart	23
Completing the cable connection table	24
Cable connection table	25
Completing the configuration data table	27
Configuration data table	27
Chapter 4. Planning guidelines for using your SAN Volume Controller in a SAN environment.	29
Storage Area Network	29
Switch zoning for the SAN Volume Controller	30
Zoning considerations for Metro Mirror	33
Switch operations over long distances	34
Operation of a SAN Volume Controller cluster in SAN fabrics with long distance fibre links	35
Performance of fibre-channel extenders	36
Nodes	36

	Clusters	37
	Cluster state	38
	Cluster operation and quorum disks	38
	I/O groups and uninterruptible power supply	39
	Uninterruptible power supply and power domains	40
	Disk controllers	42
	Data migration	43
	Image mode virtual disk migration	43
I	Copy Services	43
I	FlashCopy	44
I	FlashCopy mappings.	44
I	FlashCopy consistency groups	46
I	Metro Mirror	48
I	IBM TotalStorage Metro Mirror for SAN Volume Controller	49
I	Metro Mirror consistency groups	49
	Chapter 5. Object descriptions	51
	Storage subsystems	52
	Managed disks	54
	Managed disk groups	56
	Virtual disks	59
	Virtual disk-to-host mapping	61
	Host objects	63
	Chapter 6. Planning for configuring the SAN Volume Controller	65
	Maximum configuration	66
	Configuration rules and requirements.	68
	Configuration rules	69
	Storage subsystems	70
	Host bus adapters.	74
	Nodes	75
	Power requirements	76
	Fibre-channel switches	76
	Configuration requirements	80
	Chapter 7. SAN Volume Controller supported environment	83
	Supported host attachments	83
	Supported storage subsystems	83
	Supported fibre-channel host bus adapters	83
	Supported switches	84
	Supported fibre-channel extenders.	84
	Accessibility	85
	Notices	87
	Trademarks	88
	Definitions of notices.	88
	Glossary	91
	Index	99

Figures

1.	Levels of virtualization	2
2.	Symmetrical virtualization	4
3.	A SAN Volume Controller node	6
4.	Example of a SAN Volume Controller in a fabric	7
5.	Uninterruptible power supply 5125	9
6.	Uninterruptible power supply 5115	9
7.	I/O groups and uninterruptible power supply 5125 relationship	11
8.	Cluster, nodes, and cluster state.	38
9.	I/O group and uninterruptible power supply	40
10.	Relationship between I/O groups and uninterruptible power supply units	41
11.	Objects in a virtualized system.	52
12.	Controllers and MDisks	55
13.	MDisk group	56
14.	Managed disk groups and VDIs	60
15.	Hosts, WWPNs, and VDIs	63
16.	Hosts, WWPNs, VDIs and SCSI mappings	63
17.	Disk controller system shared between SAN Volume Controller and a host	72
18.	ESS LUs accessed directly with a SAN Volume Controller	73
19.	FASTT direct connection with a SAN Volume Controller on one host	74
20.	Fabric with Inter-Switch Links between nodes in a cluster	79
21.	Fabric with Inter-Switch Links in a redundant configuration	79

Tables

I	1. Uninterruptible power supply support guidelines	10
	2. Sample of completed hardware location chart	22
	3. Hardware location chart	23
	4. Cable connection table	25
	5. Example of cable connection table	26
	6. Four hosts and their ports	31
	7. Six hosts and their ports	32
	8. Node state	36
I	9. Required uninterruptible power supply (UPS) units	41
	10. Managed disk status	55
	11. Managed disk group status	57
I	12. Capacities of the cluster given extent size	59
	13. Virtual disk status	61
	14. SAN Volume Controller maximum configuration values	66
	15. Configuration Terms and Definitions	68

About this guide

This publication introduces the IBM® TotalStorage® SAN Volume Controller, its components, and its features.

It also provides planning guidelines for installing and configuring the SAN Volume Controller.

Who should use this guide?

This publication is intended for anyone who is planning to install and configure an IBM TotalStorage SAN Volume Controller.

Summary of changes

This document contains terminology, maintenance, and editorial changes.

Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. This summary of changes describes new functions that have been added to this release.

Summary of changes for GA22-1052-04 SAN Volume Controller Planning Guide

The Summary of Changes provides a list of new, modified, and changed information since the last version of the guide.

New information

This topic describes the changes to this guide since the previous edition, GA22-1052-03. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- The SAN Volume Controller can be used with an uninterruptible power supply (UPS) 5115. This release includes documentation on the UPS 5115 as well as the UPS 5125.

Note: If text is referring to the “UPS” or to the “uninterruptible power supply,” then it is referring to a generic UPS and can be referring to either UPS. When the UPS is referred to as the “UPS 5115” or the “UPS 5125,” then the specific UPS is designated.

- Added the following new topics:
 - Operation of the SAN Volume Controller cluster in SAN fabrics with long distance fibre links
 - Power cables for the uninterruptible power supply 5115
 - Power cables for the uninterruptible power supply 5125

Changed information

This section lists the updates that were made in this document.

- Synchronous Remote Copy and Peer-to-Peer Remote Copy (PPRC) have been renamed to Metro Mirror.

- Support for 4-node configurations has been updated to support 8-node.
- Updated support information for split controller configurations.
- Updated the number of uninterruptible power supplies to support up to 8 nodes.
- Updated the maximum configurations table.

Deleted information

Documentation for the Master Console has been removed. You can find information for the Master Console in *IBM TotalStorage Master Console Installation and User's Guide*.

The following topics were deleted for the Master Console:

- Master console
- Master console components
- Preparing your master console environment

Summary of changes for GA22-1052-03 SAN Volume Controller Planning Guide

The Summary of Changes provides a list of new, modified, and changed information since the last version of the guide.

New information

This topic describes the changes to this guide since the previous edition, GA22-1052-02. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Clusters can contain from one to four pairs of nodes.
- Software which is installed on the master console is listed.
- A cluster must have two to four uninterruptible power supply units depending on the number of nodes.
- Except when migrating between groups, a VDisk must be associated with just one MDisk group.
- An MDisk can be associated with just one MDisk group.
- There is a new procedure for installing and configuring switches to create zones.

Changed information

No information has been changed in this version.

Deleted information

No deletions were made in this version.

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.
<i>Italics</i>	Text in <i>italics</i> 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.

SAN Volume Controller library and related publications

A list of other publications that are related to this product are provided to you for your reference.

The tables in this section list and describe the following publications:

- The publications that make up the library for the IBM TotalStorage SAN Volume Controller
- Other IBM publications that relate to the SAN Volume Controller

SAN Volume Controller library

The following table lists and describes the publications that make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the SAN Volume Controller. If you need additional copies of this CD, the order number is SK2T-8811. These publications are also available as PDF files from the following Web site:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller: CIM Agent Developer's Reference</i>	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	SC26-7590
<i>IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7544
<i>IBM TotalStorage SAN Volume Controller: Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller.	SC26-7543
<i>IBM TotalStorage SAN Volume Controller: Host Attachment Guide</i>	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7575

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller: Installation Guide</i>	This guide includes the instructions the service representative uses to install the SAN Volume Controller.	SC26-7541
<i>IBM TotalStorage SAN Volume Controller: Planning Guide</i>	This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA22-1052
<i>IBM TotalStorage SAN Volume Controller: Service Guide</i>	This guide includes the instructions the service representative uses to service the SAN Volume Controller.	SC26-7542
<i>IBM TotalStorage SAN Volume Controller: Translated Safety Notices</i>	This guide contains the danger and caution notices for the SAN Volume Controller. The notices are shown in English and in numerous other languages.	SC26-7577
<i>IBM TotalStorage Master Console Installation and User's Guide</i>	This guide includes the instructions on how to install and use the SAN Volume Controller Console	

Other IBM publications

The following table lists and describes other IBM publications that contain additional information related to the SAN Volume Controller.

Title	Description	Order number
<i>IBM TotalStorage Enterprise Storage Server, IBM TotalStorage SAN Volume Controller, IBM TotalStorage SAN Volume Controller for Cisco MDS 9000, IBM TotalStorage Multipath Subsystem Device Driver: User's Guide</i>	This guide describes the IBM TotalStorage Multipath Subsystem Device Driver Version 1.5 for TotalStorage Products and how to use it with the SAN Volume Controller. This publication is referred to as the <i>IBM TotalStorage Multipath Subsystem Device Driver: User's Guide</i> .	SC26-7608

Related Web sites

The following Web sites provide information about the SAN Volume Controller or related products or technologies.

Type of information	Web site

SAN Volume Controller support	http://www-1.ibm.com/servers/storage/support/virtual/2145.html
Technical support for IBM storage products	http://www.ibm.com/storage/support/

How to order IBM publications

The publications center is a worldwide central repository for IBM product publications and marketing material.

The IBM publications center

The IBM publications center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download free of charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM publications center through the following Web site:

www.ibm.com/shop/publications/order/

Publications notification system

The IBM publications center Web site offers you a notification system for IBM publications. Register and you can create your own profile of publications that interest you. The publications notification system sends you a daily e-mail that contains information about new or revised publications that are based on your profile.

If you want to subscribe, you can access the publications notification system from the IBM publications center at the following Web site:

www.ibm.com/shop/publications/order/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this book or any other documentation, you can submit them in one of the following ways:

- e-mail

Submit your comments electronically to the following e-mail address:

starpubs@us.ibm.com

Be sure to include the name and order number of the book and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail

Fill out the Readers' Comments form (RCF) at the back of this book. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing Department
Department 61C

9032 South Rita Road
Tucson, Arizona 85775-4401
U.S.A.

Chapter 1. Virtualization and the IBM TotalStorage SAN Volume Controller

Before you begin planning the installation, understand why virtualization is needed, what virtualization is, and the IBM TotalStorage SAN Volume Controller system.

Virtualization

Virtualization is a concept that applies to many areas of the information technology industry.

For data storage, virtualization includes the creation of a pool of storage that contains several disk subsystems. These subsystems can be supplied from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them. Therefore, virtual disks can use mixed back-end storage and provide a common way to manage a storage area network (SAN).

Historically, the term *virtual storage* has described the virtual memory techniques that have been used in operating systems. The term *storage virtualization*, however, describes the shift from managing physical volumes of data to logical volumes of data. This shift can be made on several levels of the components of storage networks. Virtualization separates the representation of storage between the operating system and its users from the actual physical storage components. This technique has been used in mainframe computers for many years through methods such as system-managed storage and products like the IBM Data Facility Storage Management Subsystem (DFSMS). Virtualization can be applied at four main levels:

- Virtualization at the *server* level is performed by managing volumes on the operating systems servers. An increase in the amount of logical storage over physical storage is suitable for environments that do not have storage networks.
- Virtualization at the *storage device* level is in common use. Striping, mirroring, and redundant arrays of independent disks (RAIDs) are used by almost all disk subsystems. This type of virtualization can range from simple RAID controllers to advanced volume management such as that provided by the IBM TotalStorage Enterprise Storage Server (ESS) or by Log Structured Arrays (LSA). The Virtual Tape Server (VTS) is another example of virtualization at the device level.
- Virtualization at the *fabric* level enables storage pools to be independent of the servers and the physical components that make up the storage pools. One management interface can be used to manage different storage systems without affecting the servers. The SAN Volume Controller performs virtualization at the fabric level.
- Virtualization at the *file system* level provides the highest benefit because data is shared, allocated, and protected, not volumes.

Virtualization is a radical departure from traditional storage management. In traditional storage management, storage is attached directly to a host system, which controls storage management. SANs introduced the principle of networks of storage, but storage is still primarily created and maintained at the RAID subsystem level. Multiple RAID controllers of different types require knowledge of, and software that is specific to, the given hardware. Virtualization brings a central point of control for disk creation and maintenance. It brings new ways of managing storage maintenance.

Where storage is concerned, one problematic area that virtualization addresses is unused capacity. Rather than individual storage systems remaining islands unto themselves, allowing excess storage capacity to be wasted when jobs do not require it, storage is pooled so that jobs that need the highest storage capacity can use it when they need it. Regulating the amount of available storage becomes easier to orchestrate without computing resource or storage resource having to be turned off and on.

Types of virtualization

Virtualization can be performed either asymmetrically or symmetrically. Figure 1 provides a diagram of the levels of virtualization.

Asymmetric

A virtualization engine is outside the data path and performs a metadata style service.

Symmetric

A virtualization engine sits in the data path, presenting disks to the hosts but hiding the physical storage from the hosts. Advanced functions, such as cache and Copy Services, can therefore be implemented in the engine itself.

Virtualization at any level provides benefits. When several levels are combined, however, the benefits of those levels can also be combined. For example, you can gain the most benefits if you attach a low cost RAID controller to a virtualization engine that provides virtual volumes for use by a virtual file system.

Note: The SAN Volume Controller implements fabric-level *virtualization*. Within the context of the SAN Volume Controller and throughout this document, *virtualization* refers to symmetric fabric-level virtualization.




Figure 1. Levels of virtualization

Related concepts

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

The need for virtualization

Storage is a facility that computer users want to access at any time, from any location, with a minimum amount of management.

Users expect the storage devices to provide enough capacity and to be reliable. The amount of storage that users require, however, is increasing quickly. Internet users use large amounts of storage daily. Many users are mobile, access patterns cannot be predicted, and the content of the data becomes more and more interactive. Because the amount of data that processed is large, it can no longer be managed manually. Automatic management is required, as are new levels of bandwidth and load balancing. Also, it is important that all this data can be shared between different types of computer platforms, because the communication networks cannot process the large replication, download, and copying operations that are required.

Storage area networks (SANs) are high speed switched networks that let multiple computers share access to many storage devices. SANs allow for the use of advanced software that automatically manages the storage of data. With such advanced software, the computers that are connected to a particular network can, therefore, access storage wherever that storage is available in the network. The user is no longer aware of, and no longer needs to know, which physical devices contain which data. The storage has become virtualized. In a similar way to how virtual memory has solved the problems of the management of a limited resource in application programs, the virtualization of storage has given users a more intuitive use of storage, while software quietly manages the storage network in the background.

Fabric level virtualization models

In traditional storage management, storage devices are connected directly to host systems and are maintained locally by those host systems.

Although storage area networks (SANs) have introduced the principle of networks, storage devices are still mainly assigned to individual host systems and storage is still mainly created and maintained at the RAID subsystem level. Therefore, RAID controllers of different types need knowledge of, and software that is specific to, the hardware that is used.

Virtualization provides a complete change from the traditional storage management. It provides a central point of control for disk creation and management, and therefore requires changes to the way in which storage management is done.

Fabric level virtualization is the principle in which a pool of storage is created from more than one disk subsystem. This pool is then used to set up virtual disks that are made visible to the host systems. These virtual disks use whatever storage is available and permit a common way to manage SAN storage.

Fabric level virtualization can be done in either of two ways: asymmetric or symmetric.

With asymmetric virtualization, the virtualization engine is outside the data path. It provides a metadata server that contains all the mapping and the locking tables. The storage devices contain only data.

Because the flow of control is separated from the flow of data, input/output (I/O) operations can use the full bandwidth of the SAN. A separate network or SAN link is used for control.

However, there are disadvantages to asymmetric virtualization:

- Data is at risk to increased security exposures, and the control network must be protected with a firewall.
- Metadata can become very complicated when files are distributed across several devices.
- Each host that accesses the SAN must know how to access and interpret the metadata. Specific device driver or agent software must therefore be running on each of these hosts.
- The metadata server cannot run advanced functions, such as caching or Copy Services, because it only has access to the metadata, not the data itself.

Symmetric virtualization

The SAN Volume Controller provides symmetric virtualization.

Virtualization splits the physical storage Redundant Array of Independent Disks (RAID) arrays into smaller chunks of storage that are known as extents. These extents are then concatenated, using various policies, to make virtual disks. With symmetric virtualization, host systems can be isolated from the physical storage. Advanced functions, such as data migration, can run without the need to re-configure the host. With symmetric virtualization, the virtualization engine is the central configuration point for the SAN.

In symmetric virtual storage networks (see Figure 2), data and control both flow over the same path. Because the separation of the control from the data occurs in the data path, the storage can be pooled under the control of the virtualization engine. The virtualization engine performs the logical-to-physical mapping.




Figure 2. Symmetrical virtualization

The virtualization engine directly controls access to the storage and to the data that is written to the storage. As a result, locking functions that provide data integrity and

advanced functions, such as cache and Copy Services, can be run in the virtualization engine itself. The virtualization engine is, therefore, a central point of control for device and advanced function management. Symmetric virtualization also allows you to build a kind of firewall in the storage network. Only the virtualization engine can give access through the firewall. Symmetric virtualization does, however, cause some problems.

The main problem that is associated with symmetric virtualization is related to poor performance, because all I/O must flow through the virtualization engine. This problem is one of scalability. You can use an n-way cluster of virtualization engines that has failover capacity to solve this problem. You can scale the additional processor power, cache memory, and adapter bandwidth to get the level of performance that you want. The memory and processing power can run the advanced functions, such as Copy Services and caching.

The IBM TotalStorage SAN Volume Controller uses symmetric virtualization. Single virtualization engines, which are known as nodes, are combined to create clusters. Each cluster can contain between two and eight nodes.

Related concepts

“Virtualization” on page 1

Virtualization is a concept that applies to many areas of the information technology industry.

SAN Volume Controller

The SAN Volume Controller is a SAN appliance that attaches open-systems storage devices to supported open-systems hosts.

The IBM TotalStorage SAN Volume Controller provides symmetric virtualization by creating a pool of managed disks from the attached storage subsystems, which are then mapped to a set of virtual disks for use by attached host computer systems. System administrators can view and access a common pool of storage on the SAN, which enables them to use storage resources more efficiently and provides a common base for advanced functions.

The SAN Volume Controller is analogous to a logical volume manager (LVM) on a SAN. It performs the following functions for the SAN storage that it is controlling:

- Creates a single pool of storage
- Logical unit virtualization
- Manages logical volumes
- Provides advanced functions for the SAN, such as:
 - Large scalable cache
 - Copy services
 - FlashCopy® (point-in-time copy)
 - Metro Mirror (synchronous copy)
 - Data migration
 - Space management
 - Mapping that is based on desired performance characteristics
 - Quality of service metering

A *node* is a single storage engine. Figure 3 on page 6 provides an illustration of a node. The storage engines are always installed in pairs with one to four pairs of

nodes constituting a *cluster*. Each node in a pair is configured to back up the other. Each pair of nodes is known as an *I/O group*. All I/O operations that are managed by the nodes in an I/O group are cached on both nodes for resilience. Each virtual volume is defined to an I/O group. To avoid any single point of failure, the nodes of an I/O group are protected by independent uninterruptible power supply (UPS) 5115 or uninterruptible power supply (UPS) 5125 units.




Figure 3. A SAN Volume Controller node

The SAN Volume Controller I/O groups recognize the storage presented to the SAN by the backend controllers as a number of disks known as *managed disks*. The application servers do not recognize these managed disks. Instead they see a number of logical disks, known as *virtual disks*, that are presented to the SAN by the SAN Volume Controller. Each node must be in only one I/O group and provide access to the virtual disks in the I/O group.

The SAN Volume Controller helps to provide continuous operations and can also optimize the data path to ensure performance levels are maintained. Ensure that you use IBM TotalStorage Multiple Device Manager performance manager to analyze the performance statistics. See *IBM TotalStorage Multiple Device Manager Configuration and Installation Guide* and *IBM TotalStorage Multiple Device Manager CLI Guide* for more information.

The fabric contains two distinct zones: a host zone and a disk zone. In the host zone, the host systems can identify and address the nodes. You can have more than one host zone. Generally, you will create one host zone per operating system type. In the disk zone, the nodes can identify the disk drives. Host systems cannot operate on the disk drives directly; all data transfer occurs through the nodes. Figure 4 on page 7 shows that several host systems can be connected to a SAN fabric. A cluster of SAN Volume Controller nodes is connected to the same fabric and presents virtual disks to the host systems. You create these virtual disks from units of space within a managed disk group. A managed disk group is a collection of managed disks presented by the back-end RAID controllers, providing a storage pool. You choose how each group is made up, and you can combine managed disks from different manufacturers' controllers in the same managed disk group if required.

Note: Some operating systems cannot tolerate other operating systems in the same zone.

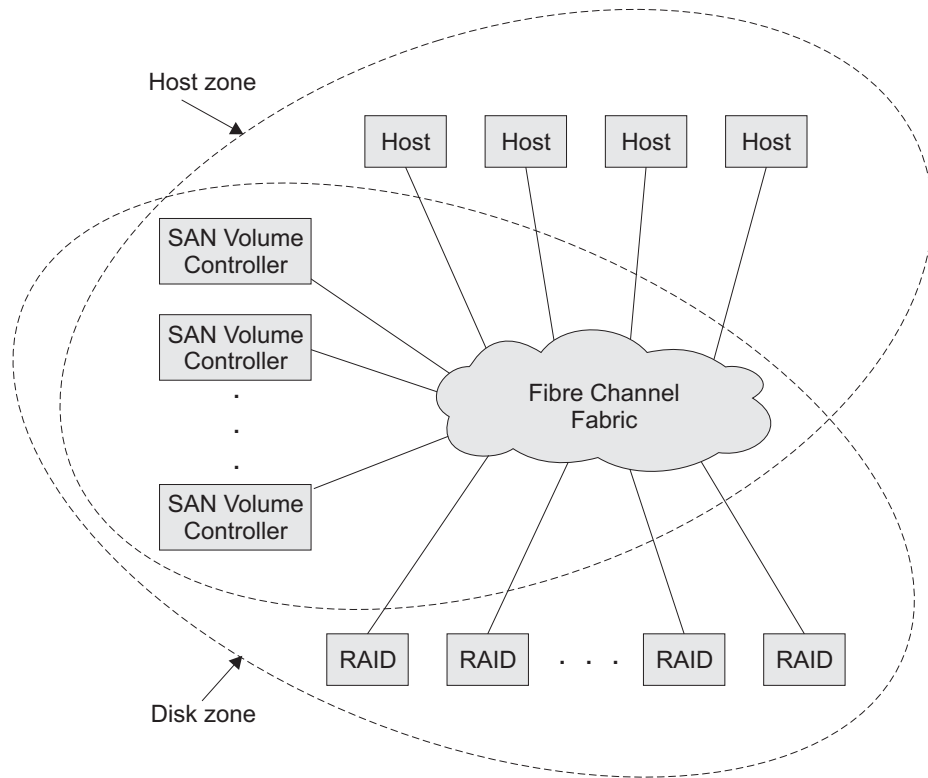


Figure 4. Example of a SAN Volume Controller in a fabric

You can remove one node in each I/O group from a cluster when hardware service or maintenance is required. After you remove the node, you can replace the field replaceable units (FRUs) in the node. All disk drive communication and communication between nodes is performed through the SAN. All SAN Volume Controller configuration and service commands are sent to the cluster through an Ethernet network.

Each node contains its own vital product data (VPD). Each cluster contains VPD that is common to all the nodes on the cluster, and any system connected to the Ethernet network can access this VPD.

Cluster configuration information is stored on every node that is in the cluster to allow concurrent replacement of FRUs. An example of this information might be information that is displayed on the menu screen of the SAN Volume Controller. When a new FRU is installed and when the node is added back into the cluster, configuration information that is required by that node is read from other nodes in the cluster.

SAN Volume Controller operating environment

- Minimum of one pair of SAN Volume Controller nodes
- Minimum of two uninterruptible power supplies
- One master console is required per SAN installation for configuration

Features of a SAN Volume Controller node

- 19-inch rack mounted enclosure
- 4 fibre channel ports
- 2 fibre channel adapters
- 4 GB cache memory

Supported hosts

For a list of supported operating systems, see the IBM TotalStorage SAN Volume Controller Web site at:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Multipathing software

- IBM Subsystem Device Driver (SDD)
- Redundant Dual Active Controller (RDAC)

Note: Directly-attached hosts that share a back-end storage controller with a SAN Volume Controller can run multipath drivers SDD and RDAC simultaneously. Other multipath drivers running with SDD are not supported.

Check the following Web site for the latest support and coexistence information:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

User interfaces

The SAN Volume Controller provides the following user interfaces:

- IBM TotalStorage SAN Volume Controller Console, a Web-accessible graphical user interface (GUI) that supports flexible and rapid access to storage management information
- A command-line interface (CLI) using Secure Shell (SSH)

Application programming interfaces

The SAN Volume Controller provides the following application programming interface:

- IBM TotalStorage Common Information Model (CIM) Agent for the SAN Volume Controller, which supports the Storage Management Initiative Specification of the Storage Network Industry Association.

Related concepts

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Related reference

“Supported host attachments” on page 83

The IBM Web site provides up-to-date information about the supported host attachment operating systems.

Uninterruptible power supply overview

The uninterruptible power supply (UPS) provides the SAN Volume Controller with a secondary power source to be used if you lose power from your primary power

|
|
|

source due to power failures, power sags, power surges, or line noise. Two types of UPS units can be used with the SAN Volume Controller: the UPS 5115 and the UPS 5125.

If a power outage occurs, the UPS maintains power long enough to save any configuration and cache data contained in the dynamic random access memory (DRAM). The data is saved to the SAN Volume Controller internal disk. Figure 5 and Figure 6 provide illustrations of the two types of UPS units.




Figure 5. Uninterruptible power supply 5125




Figure 6. Uninterruptible power supply 5115

Note: The SAN Volume Controller UPS is an integral part of the SAN Volume Controller solution. It maintains continuous SAN Volume Controller specific communications with its attached SAN Volume Controller nodes. The SAN Volume Controller cannot operate without the UPS. The SAN Volume

Controller UPS must be used in accordance with documented guidelines and procedures and must not power any equipment other than SAN Volume Controller nodes.

To provide full redundancy and concurrent maintenance, the SAN Volume Controller must be installed in pairs. Each SAN Volume Controller of a pair must be connected to a different UPS. Each UPS 5125 can support up to two SAN Volume Controller nodes. The UPS 5115 can support one SAN Volume Controller node. It is also recommended that you connect the two UPS units for the pair to different independent electrical power sources. This reduces the chance of an input power failure at both UPS units.

The UPS must be in the same rack as the nodes. Ensure that you are following the UPS guidelines that are provided in Table 1.

Table 1. Uninterruptible power supply support guidelines

Number of nodes	Number of uninterruptible power supply 5125 units	Number of uninterruptible power supply 5115 units
2	2	2
4	2	4
6	4	6
8	4	8

Attention:

1. Do not connect the uninterruptible power supplies to an input power source that does not conform to standards.
2. Each UPS pair must power only one SAN Volume Controller cluster.

Each UPS includes power (line) cords that connect the UPS to either a rack power distribution unit (PDU), if one exists, or to an external power source. Each UPS power input requires the protection of a UL-approved (or equivalent) 250 volt, 15 amp circuit breaker.

The UPS is connected to the SAN Volume Controllers with a power cable and a signal cable. To avoid the possibility of power and signal cables being connected to different UPS units, these cables are wrapped together and supplied as a single field replaceable unit. The signal cables enable the SAN Volume Controllers to read status and identification information from the UPS.

Each SAN Volume Controller monitors the operational state of the UPS to which it is attached. If the UPS reports a loss of input power, the SAN Volume Controller stops all I/O operations and dumps the contents of its DRAM to the internal disk drive. When input power to the UPS is restored, the SAN Volume Controllers restart and restore the original contents of the DRAM from the data saved on the disk drive.

A SAN Volume Controller is not fully operational until the UPS battery charge state indicates that it has sufficient capacity to power the SAN Volume Controller long enough to permit it to save all its memory to the disk drive in the event of a power loss. The UPS has sufficient capacity to save all the data on the SAN Volume Controller at least twice. For a fully-charged UPS, even after battery capacity has been used to power the SAN Volume Controllers while they save DRAM data,

sufficient battery capacity remains to let the SAN Volume Controllers become fully operational as soon as input power is restored.

Note: Under normal circumstances, if input power is disconnected from the UPS, the SAN Volume Controller connected to that UPS performs a power down sequence. This operation, which saves the configuration and cache data to an internal disk in the SAN Volume Controller, typically takes about three minutes, at which time power is removed from the output of the UPS. In the event of a delay in the completion of the power down sequence, the UPS output power is removed five minutes after the power was disconnected to the UPS. Because this operation is controlled by the SAN Volume Controller, a UPS that is not connected to an active SAN Volume Controller will not shut off within the five-minute required period. In the case of an emergency, you must manually shut down the UPS by pushing the UPS 5125 power off button, or the UPS 5115 on/off button.

Attention: Data integrity could be compromised by pushing the UPS 5125 power off button or the UPS 5115 on/off button. Never shut down a UPS without first shutting down the SAN Volume Controller nodes that it supports.

It is very important that the two nodes in the I/O group are connected to different uninterruptible power supplies. This configuration ensures that cache and cluster state information is protected in the event of a failure of the UPS or mainline power source.

When nodes are added to the cluster, you must specify the I/O group that they will join. The configuration interfaces also checks the UPS units and ensures that the two nodes in the I/O group are not connected to the same UPS units.

Figure 7 shows a cluster of four nodes, with two I/O groups and two UPS 5125 units.




Figure 7. I/O groups and uninterruptible power supply 5125 relationship

Overview of backup functions

The SAN Volume Controller includes functions that help you to back up cluster configuration settings and business data.

To enable routine maintenance of the SAN Volume Controller clusters, the configuration settings for each cluster are stored on each node. If power fails on a cluster or if a node in a cluster is replaced, the cluster configuration settings are automatically restored when the repaired node is added to the cluster. To restore the cluster configuration in the event of a disaster (if all nodes in a cluster are lost simultaneously), plan to back up the cluster configuration settings to tertiary storage. You use the configuration backup functions to back up the cluster configuration.

For complete disaster recovery, regularly back up the business data that is stored on virtual disk at the application server level or the host level. The SAN Volume Controller provides the following Copy Services functions that you can use to back up data: Metro Mirror and FlashCopy.

Related concepts

“Metro Mirror” on page 48

Metro Mirror (formerly known as Remote Copy) enables you to set up a relationship between two virtual disks, so that updates that are made by an application to one virtual disk are mirrored on the other virtual disk.

“FlashCopy” on page 44

FlashCopy is a copy service that is available with the SAN Volume Controller.

Cluster configuration backup

Configuration backup is the process of extracting configuration data from a cluster and writing it to disk.

Backing up the cluster configuration enables you to restore it in the event that configuration data is lost. The data that is backed up is the metadata that describes the cluster configuration, not the data that your enterprise uses to run its business.

The backup configuration files can be saved on the master console or the configuration node.

Objects included in the backup

Configuration data is information about a cluster and the objects that are defined in it. The following objects are copied:

- Storage subsystem
- Hosts
- I/O groups
- Managed disks (MDisks)
- MDisk groups
- Nodes
- Virtual disks (VDisks)
- VDisk-to-host mappings
- SSH key
- FlashCopy mappings
- FlashCopy consistency groups
- Metro Mirror relationships
- Metro Mirror consistency groups

Related concepts

“Clusters” on page 37
All configuration and service is performed at the cluster level.

Chapter 2. Installation planning

Before the service representative can start to set up your SAN Volume Controller, you must verify that the prerequisite conditions for the SAN Volume Controller and uninterruptible power supply (UPS) installation are met.

1. Does your physical site meet the environment requirements for the SAN Volume Controller and UPS?
2. Do you have adequate rack space for your hardware?
 - a. SAN Volume Controller: One Electrical Industries Association (EIA) unit high for each node.
 - b. UPS 5125: Two EIA units high for each UPS 5125.
 - c. UPS 5115: One EIA unit high for each UPS 5115.
3. Do you have power distribution units in the rack to provide power to the UPS units?

A clearly visible and accessible emergency power off switch is required.
4. Ensure that you provide appropriate connectivity.

Preparing your SAN Volume Controller environment

Before installing the SAN Volume Controller, prepare the physical environment.

Dimensions and weight

The following tables list the physical dimensions and weight of the SAN Volume Controller, as well as other environmental requirements that you must consider before you install your SAN Volume Controller:

Height	Width	Depth	Maximum Weight
43 mm (1.7 in.)	440 mm (17.3 in.)	660 mm (26 in.)	12.7 kg (28 lb.)

Additional space requirements

Location	Additional Space Required	Reason
Left and right sides	50 mm (2 in.)	Cooling air flow
Back	minimum: 100 mm (4 in.)	Cable exit

AC input-voltage requirements

Power Supply Assembly Type	Voltage	Frequency
200 to 240 V	88 to 264 V ac	50 to 60 Hz

Environment

Environment	Temperature	Altitude	Relative Humidity	Maximum Wet Bulb Temperature
Operating in Lower Altitudes	10°C to 35°C (50°F to 95°F)	0 to 914 m (0 to 2998 ft.)	8% to 80% noncondensing	23°C (74°F)
Operating in Higher Altitudes	10°C to 32°C (50°F to 88°F)	914 to 2133 m (2998 to 6988 ft.)	8% to 80% noncondensing	23°C (74°F)
Powered Off	10°C to 43°C (50°F to 110°F)	–	8% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 6988 ft.)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft.)	5% to 100% condensing, but no precipitation	29°C (84°F)

Heat output (maximum)

350 watts (1195 Btu per hour)

Power cables for the uninterruptible power supply 5115

You must follow your country or region's power requirements to choose the appropriate power cable for the uninterruptible power supply 5115.

The following table lists the power cable requirements for your country or region.

Country or region	Length	Connection type (attached plug designed for 200-240V AC input)	Part
United States of America (Chicago), Canada, Mexico	1.8 m (6 ft)	NEMA L6-15P	7842122
Bahamas, Barbados, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Korea (South), Liberia, Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Saudi Arabia, Suriname, Taiwan, Trinidad (West Indies), United States of America, Venezuela	2.8 m (9 ft)	NEMA L6-15P	7842123

Country or region	Length	Connection type (attached plug designed for 200-240V AC input)	Part
Antigua, Bahrain, Brunei, Channel Islands, China (Hong Kong S.A.R.), Cyprus, Dubai, Fiji, Ghana, India, Iraq, Ireland, Kenya, Kuwait, Malawi, Malaysia, Malta, Nepal, Nigeria, Polynesia, Qatar, Sierra Leone, Singapore, Tanzania, Uganda, United Kingdom, Yemen, Zambia	2.8 m (9 ft)	BS 1363/A	14F0033
Argentina, Australia, China (PRC), New Zealand, Papua New Guinea, Paraguay, Uruguay, Western Samoa	2.8 m (9 ft)	AZ/NZS C112	13F9940
Afghanistan, Algeria, Andorra, Angola, Austria, Belgium, Benin, Bulgaria, Burkina Faso, Burundi, Cameroon, Central African Republic, Chad, China (Macau S.A.R.), Czech Republic, Egypt, Finland, France, French Guiana, Germany, Greece, Guinea, Hungary, Iceland, Indonesia, Iran, Ivory Coast, Jordan, Lebanon, Luxembourg, Malagasy, Mali, Martinique, Mauritania, Mauritius, Monaco, Morocco, Mozambique, Netherlands, New Caledonia, Niger, Norway, Poland, Portugal, Romania, Senegal, Slovakia, Spain, Sudan, Sweden, Syria, Togo, Tunisia, Turkey, former USSR, Vietnam, former Yugoslavia, Zaire, Zimbabwe	2.8 m (9 ft)	CEE 7-VII	13F9979
Denmark	2.8 m (9 ft)	DK2-5a	13F9997
Bangladesh, Burma, Pakistan, South Africa, Sri Lanka	2.8 m (9 ft)	SABS 164	14F0015
Liechtenstein, Switzerland	2.8 m (9 ft)	1011-S2450 7	14F0051
Chile, Ethiopia, Italy, Libya, Somalia	2.8 m (9 ft)	CEI 23-16	14F0069
Israel	2.8 m (9 ft)	SI 32	14F0087

Power cables for the uninterruptible power supply 5125

You must follow your country or region's power requirements to choose the appropriate power cable for the uninterruptible power supply 5125.

The following table lists the power cable requirements for your country or region.

Country or region	Length	Connection type (attached plug designed for 200-240V AC input)	Part
United States of America (Chicago), Canada, Mexico	1.8 m (6 ft)	NEMA L6-15P	14F1549
Bahamas, Barbados, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Korea (South), Liberia, Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Saudi Arabia, Suriname, Taiwan, Trinidad (West Indies), United States of America, Venezuela	2.5 m (8 ft)	NEMA L6-15P	12J5119
Antigua, Bahrain, Brunei, Channel Islands, China (Hong Kong S.A.R.), Cyprus, Denmark, Dubai, Fiji, Ghana, India, Iraq, Ireland, Kenya, Kuwait, Malawi, Malaysia, Malta, Nepal, Nigeria, Polynesia, Qatar, Sierra Leone, Singapore, Tanzania, Uganda, United Kingdom, Yemen, Zambia	2.5 m (8 ft)	IEC 309	36L8822
Argentina, Australia, China (PRC), New Zealand, Papua New Guinea, Paraguay, Uruguay, Western Samoa	2.5 m (8 ft)	L6-20P	12J5118
Afghanistan, Albania, Algeria, Andorra, Angola, Austria, Belgium, Benin, Bulgaria, Burkina Faso, Burundi, Cameroon, Central African Republic, Chad, China (Macau S.A.R.), Czech Republic, Egypt, Finland, France, French Guiana, Germany, Greece, Guinea, Hungary, Iceland, Indonesia, Iran, Ivory Coast, Jordan, Lebanon, Luxembourg, Malagasy, Mali, Martinique, Mauritania, Mauritius, Monaco, Morocco, Mozambique, Netherlands, New Caledonia, Niger, Norway, Poland, Portugal, Romania, Senegal, Slovakia, Spain, Sudan, Sweden, Syria, Togo, Tunisia, Turkey, former USSR, Vietnam, former Yugoslavia, Zaire, Zimbabwe	2.5 m (8 ft)	CEE7	55H6643
Bangladesh, Burma, Pakistan, South Africa, Sri Lanka	2.5 m (8 ft)	SABS 164	12J5124
Thailand	2.5 m (8 ft)	NEMA 6-15P	12J5120

Preparing your uninterruptible power supply environment

Ensure that your physical site meets the installation requirements for the uninterruptible power supply (UPS).

Use the following considerations when you configure the UPS 5115:

- The voltage that is supplied to the UPS 5115 must be 220 – 240 V, single phase (note that the UPS 5115 has an integrated circuit breaker and does not require external protection).

Use the following considerations when configuring the UPS 5125:

- Each UPS 5125 must be connected to a separate branch circuit.
- A UL listed 15 A circuit breaker must be installed in each branch circuit that supplies power to the UPS 5125.
- The voltage that is supplied to the UPS 5125 must be 200 – 240 V, single phase.
- The frequency supplied must be between 50 and 60 Hz.

Attention: Ensure that you comply with the following requirements for uninterruptible power supplies.

Note: If the UPS is cascaded from another UPS, the source UPS must have at least three times the capacity per phase, and the total harmonic distortion must be less than 5% with any single harmonic being less than 1%. The UPS must also have input voltage capture that has a slew rate faster than 3 Hz per second and 1 msec glitch rejection.

Uninterruptible Power Supply 5115 Dimensions and Weight

Height	Width	Depth	Maximum weight
44 mm (1.73 in.)	439 mm (17.3 in.)	579 mm (22.8 in.)	18.8 kg (41.4 lb.) Packaged

Uninterruptible Power Supply 5125 Dimensions and Weight

Height	Width	Depth	Maximum weight
89 mm (3.5 in.)	483 mm (19 in.)	622 mm (24.5 in.)	37 kg (84 lb.)

AC Input-Voltage Requirements

	UPS 5115	UPS 5125
Power Rating	750 VA/520 W	3000 VA/2700 W
Voltage	220/230/240 V	200 – 240 V
Frequency	50 – 60 Hz	50 – 60 Hz

Environment

	Operating Environment	Non-operating Environment	Storing Environment	Shipping Environment
Air Temperature	0°C – 40°C (32°F – 104°F)	0°C – 40°C (32°F – 104°F)	0°C – 25°C (32°F – 77°F)	–25°C – 55°C (–13°F – 131°F)
Relative Humidity	5% – 95% non-condensing	5% – 95% non-condensing	5% – 95% non-condensing	5% – 95% non-condensing

Altitude

	Operating Environment	Non-operating Environment	Storing Environment	Shipping Environment
Altitude (from sea level)	0 – 2000 m (0 – 6560 ft.)	0 – 2000 m (0 – 6560 ft.)	0 – 2000 m (0 – 6560 ft.)	0 – 15 000 m (0 – 49212 ft.)

Heat output (maximum)

142 watts (485 Btu per hour) during normal operation.

553 watts (1887 Btu per hour) when power has failed and the UPS is supplying power to the nodes of the SAN Volume Controller.

Ports and connections

Each SAN Volume Controller requires the following ports and connections:

- Each SAN Volume Controller node requires one Ethernet cable to connect it to an Ethernet switch or hub. A 10/100 Mb Ethernet connection is required.
- Two TCP/IP addresses are normally required for a SAN Volume Controller cluster, a cluster address and a service address.
- Each SAN Volume Controller node has four fibre-channel ports, which are supplied fitted with LC-style optical Small Form-factor Pluggable (SFP) GBICs for connection to a fibre-channel switch.

Each uninterruptible power supply requires the following cable type:

- Serial cables that connect the uninterruptible power supply to the SAN Volume Controller nodes. Ensure that for each node, the serial and power cables come from the same uninterruptible power supply.

Chapter 3. Preparing the physical configuration

Before the service representative installs the SAN Volume Controller, uninterruptible power supply unit, and master console, you must plan the physical configuration and the initial settings for the system.

To plan the configuration, print or photocopy the blank charts and tables in this publication and use a pencil or pen to plan the system configuration. Before you begin writing in the charts and tables, make copies of the blank charts and tables so that you can later revise the configuration or create a new one if needed.

1. Use the hardware location chart to record the physical configuration of your system.
2. Use the cable connection table to record how your SAN Volume Controller, the uninterruptible power supply unit, and the master console are to be connected.
3. Use the configuration data table to record the data that you and the service representative need before the initial installation.

When you have completed these tasks you are ready to perform the physical installation.

Completing the hardware location chart

The hardware location chart represents the rack into which the SAN Volume Controller is installed. Each row of the chart represents one Electrical Industries Association (EIA) 19-inch rack space.

- Uninterruptible power supply units are heavy. Install them as near the bottom of the rack as possible. IBM recommends that you place them within the range of row 1 through row 8.
- Do not exceed the maximum power rating of the rack and input power supply.
- Position the SAN Volume Controller such that information on the display screen is easily viewed and the controls used to navigate the display menu are easily reached. The recommended range is EIA 11-38.
- To enable easy access to the connectors on the rear of the master console, position the console, keyboard and monitor unit adjacent to each other. To permit easy access to the CD drive, locate the master console above the keyboard and monitor unit. The recommended range is EIA 17-24.
- A SAN Volume Controller is one EIA unit high. Therefore, for each SAN Volume Controller that is to be installed, fill in the row that represents the position that the SAN Volume Controller is to occupy.
- The uninterruptible power supply 5125 is two EIA units high. Therefore, for each uninterruptible power supply 5125, fill in two rows.
- The uninterruptible power supply 5115 is one EIA unit high. Therefore, for each uninterruptible power supply 5115, fill in one row.
- The master console is two EIA units high: one EIA unit for the server and one EIA unit for the keyboard and monitor.
- If there are any hardware devices already in the rack, record this information on the chart.
- Fill in rows for all other units that will be present in the rack, including Ethernet hubs and fibre-channel switches. Hubs and switches are usually one EIA unit high, but check with your supplier. The uninterruptible power supply units must

be installed at the bottom of the rack because it might be necessary to relocate some other devices before the SAN Volume Controller installation is started.

Hardware location guidelines

When you fill in the hardware location chart, follow these basic guidelines.

- Install the SAN Volume Controller in pairs to provide redundancy and concurrent maintenance.
- A cluster can contain no more than four SAN Volume Controllers.
- Connect each SAN Volume Controller of a pair to a different uninterruptible power supply.
- Each uninterruptible power supply pair can support one SAN Volume Controller cluster.
- To reduce the chance of a simultaneous input power failure at both uninterruptible power supply units, connect each uninterruptible power supply to a separate electrical power source on a separate branch circuit.
- Because uninterruptible power supply units are heavy, install them in the lowest available positions of the rack. If necessary, move any lighter units that are already in the rack to higher positions.
- IBM does not install the Ethernet hub or the fibre-channel switches. You must arrange for either the suppliers or someone in your organization to install those items. Provide the installer with a copy of the completed hardware location chart.

In the following example, assume that the rack is empty and you want to create a system that contains the following components:

- Four SAN Volume Controllers named SVC1, SVC2, SVC3 and SVC4.
- One master console.
- Two uninterruptible power supply units named uninterruptible power supply 1 and uninterruptible power supply 2.
- One Ethernet hub named Ethernet hub 1. For this example, it is assumed that the hub is one EIA unit high.
- Two fibre-channel switches named FC switch 1 and FC switch 2. In this example, each switch is one EIA unit high.
- RAID controllers named RAID controller 1, RAID controller 2, RAID controller 3, RAID controller 4.

Your completed chart might look like Table 2:

Table 2. Sample of completed hardware location chart

Rack row	Component
EIA 36	Blank
EIA 35	Ethernet Hub 1
EIA 34	Blank
EIA 33	Blank
EIA 32	Blank
EIA 31	Blank
EIA 30	Blank
EIA 29	Blank
EIA 28	FC Switch 1

Table 2. Sample of completed hardware location chart (continued)

Rack row	Component
EIA 27	FC Switch 2
EIA 26	Blank
EIA 25	Blank
EIA 24	Blank
EIA 23	Blank
EIA 22	SAN Volume Controller 4
EIA 21	SAN Volume Controller 3
EIA 20	SAN Volume Controller 2
EIA 19	SAN Volume Controller 1
EIA 18	Master console
EIA 17	Master console keyboard and monitor
EIA 16	RAID Controller 4
EIA 15	
EIA 14	
EIA 13	RAID Controller 3
EIA 12	
EIA 11	
EIA 10	RAID Controller 2
EIA 9	
EIA 8	
EIA 7	RAID Controller 1
EIA 6	
EIA 5	
EIA 4	uninterruptible power supply 2
EIA 3	
EIA 2	uninterruptible power supply 1
EIA 1	

You might want to put the switches between the SAN Volume Controller nodes. Remember, however, that the uninterruptible power supply units must be in the lowest positions of the rack.

Hardware location chart

The hardware location chart helps you to plan the location of the hardware.

Each row of the chart in Table 3 represents one EIA unit.

Table 3. Hardware location chart

Rack row	Component
EIA 36	
EIA 35	

Table 3. Hardware location chart (continued)

Rack row	Component
EIA 34	
EIA 33	
EIA 32	
EIA 31	
EIA 30	
EIA 29	
EIA 28	
EIA 27	
EIA 26	
EIA 25	
EIA 24	
EIA 23	
EIA 22	
EIA 21	
EIA 20	
EIA 19	
EIA 18	
EIA 17	
EIA 16	
EIA 15	
EIA 14	
EIA 13	
EIA 12	
EIA 11	
EIA 10	
EIA 9	
EIA 8	
EIA 7	
EIA 6	
EIA 5	
EIA 4	
EIA 3	
EIA 2	
EIA 1	

Completing the cable connection table

The cable connection table helps you to plan how to connect the units that will be placed in the rack.

The following terms and descriptions assists you in completing the cable connection table:

Term	Description
Node number	The nominal number (name) of the SAN Volume Controller.
Uninterruptible power supply	The uninterruptible power supply to which the SAN Volume Controller is connected.
Ethernet	The Ethernet hub or switch to which the SAN Volume Controller is connected.
FC Ports 1 through 4	The fibre-channel switch ports to which the four SAN Volume Controller fibre-channel ports are connected. When viewed from the back of the SAN Volume Controller, the ports are numbered 1 through 4, from left to right. Ignore the markings on the back of the SAN Volume Controller.

For the master console, use the following terms and descriptions to complete the cable connection table:

Term	Description
Ethernet Port 1	Ethernet port 1 is used for your VPN connection. This port is required if you configure your master console to enable remote support. A remote support connection can only be enabled when this port has access to an external internet connection. For added security, you can disconnect this port when a remote support connection is not being used.
Ethernet Port 2	Ethernet port 2 is used to connect the SAN Volume Controller to the network.
FC Ports 1 and 2	FC Ports 1 and 2 are the fibre-channel switch ports to which the master console fibre-channel ports are connected. Connect one FC port to each of the SAN Volume Controller fabrics.

Cable connection table

The cable connection table helps plan the placement of units in the rack.

Complete Table 4.

Table 4. Cable connection table

SAN Volume Controller	Uninterruptible power supply	Ethernet hub or switch	FC port-1	FC port-2	FC port-3	FC port-4

Master console	Ethernet		FC port-1	FC port-2
	Public network	VPN		

Example of a completed cable connection table:

For this example, assume you are completing the cabling details for the system. Remember that SAN Volume Controllers are configured in pairs, and that the two SAN Volume Controllers of a pair must *not* be connected to the same uninterruptible power supply. Also, the two uninterruptible power supply units of a pair should not be connected to the same power source to reduce the chance of input power failure at both uninterruptible power supply units. For this example, assume that the pairs of SAN Volume Controllers are node 1 and node 2, and node 3 and node 4, and that the two power sources provided by the uninterruptible power supply units are A and B.

Note: The uninterruptible power supply requires two dedicated branch circuits that meet the following specifications:

- 15 amp circuit breaker in each branch circuit that supplies the power to an uninterruptible power supply
- Single-phase
- 50 – 60 Hz
- 220 volt

For the Ethernet connection, you must use Ethernet port 1 of the SAN Volume Controller. Do not use any other Ethernet port, because the software is configured for Ethernet port 1 only.

Note: Connect all SAN Volume Controller nodes that are part of the same cluster to the same Ethernet subnet. Otherwise, TCP/IP address failover does not work.

Table 5 illustrates this example.

Table 5. Example of cable connection table

SAN Volume Controller	Uninterruptible power supply	Ethernet hub or switch	FC Port-1	FC Port-2	FC Port-3	FC Port-4
Node 1	Uninterruptible power supply A	Hub or switch 1, Port 1	FC switch 1, Port 1	FC switch 2, Port 1	FC switch 1, Port 2	FC switch 2, Port 2
Node 2	Uninterruptible power supply B	Hub or switch 1, Port 2	FC switch 1, Port 3	FC switch 2, Port 3	FC switch 1, Port 4	FC switch 2, Port 4
Node 3	Uninterruptible power supply A	Hub or switch 1, Port 3	FC switch 1, Port 5	FC switch 2, Port 5	FC switch 1, Port 6	FC switch 2, Port 6
Node 4	Uninterruptible power supply B	Hub or switch 1, Port 4	FC switch 1, Port 7	FC switch 2, Port 7	FC switch 1, Port 8	FC switch 2, Port 8

Master console	Ethernet		FC switch 1, Port-9	FC switch 2, Port-9
	Public network	VPN		
Master console	Ethernet hub 1, Port 5	Ethernet hub 1, Port 6	FC Port-1 FC switch 1, Port 9	FC Port-2 FC switch 2, Port 9

Completing the configuration data table

The configuration data table helps you to plan the initial settings for the cluster configuration.

Include the following initial settings for the cluster:

Term	Description
Language	The national language in which you want the messages displayed on the front panel. This option applies only to service messages. The default is English.
Cluster IP address	The address that will be used for all typical configuration and service access to the cluster.
Service IP address	The address that will be used for emergency access to the cluster.
Gateway IP address	The IP address for the default local gateway for the cluster.
Subnet mask	The subnet mask of the cluster.
Fibre-channel switch speed	The fibre-channel switch speed can be either 1 Gb or 2 Gb.

Include the following information for the master console:

Term	Description
Machine name	The name you want the master console to be known as. This must be a fully qualified DNS name. The default setting is <i>mannode</i> (not fully qualified).
Master console IP addresses	The addresses that will be used for access to the master console. The default settings are: Port 1 = 192.168.1.11 Port 2 = 192.168.1.2
Master console gateway IP address	The IP address for the local gateway for the master console. The default setting is 192.168.1.1.
Master console subnet mask	The default subnet mask for the master console is 255.255.255.0.

Configuration data table

Use the configuration data table to plan the initial settings for the cluster configuration.

Cluster	
Language	
Cluster IP address	
Service IP address	
Gateway IP address	
Subnet mask	
Fibre-channel switch speed	
Master console	
Machine name	

Cluster		
	Ethernet Port 1	Ethernet Port 2
Master console IP address		
Master console gateway IP address		
Master console subnet mask		

Chapter 4. Planning guidelines for using your SAN Volume Controller in a SAN environment

Follow these planning steps to set up your SAN Volume Controller environment.

1. Plan your configuration.
2. Plan your SAN environment.
3. Plan your fabric setup.
4. Create the RAID resources that you intend to virtualize.
5. Determine if you have a RAID array that contains data that you want to merge into the cluster.
6. Determine if you will migrate data into the cluster or keep them as image-mode VDisks.
7. Determine if you will use Copy Services. Copy Services are provided for all supported hosts that are connected to the SAN Volume Controller.

Related concepts

“Image mode virtual disk migration” on page 43

Image mode virtual disks (VDisks) have a special property that allows the last extent in the VDisk to be a partial extent.

Storage Area Network

A storage area network (SAN) is a high-speed dedicated network for sharing storage resources.

A SAN allows the establishment of direct connections between storage devices and servers. It offers simplified storage management, scalability, flexibility, availability, and improved data access, movement, and backup.

A SAN storage system consists of two to eight SAN Volume Controller nodes that are arranged in a cluster. These will appear as part of the SAN fabric, along with the host systems, the RAID controllers, and the storage devices, all connected together to create the SAN. Other devices such as fabric switches might be required to complete the SAN.

It is important to note these two types of SAN: redundant and counterpart. A *redundant* SAN consists of a fault tolerant arrangement of two counterpart SANs. A redundant SAN configuration provides two independent paths for each device attached to the SAN. A *counterpart* SAN is a nonredundant portion of a redundant SAN and provides all the connectivity of the redundant SAN, but without the redundancy. Each counterpart SAN provides an alternate path for each device attached to the SAN.

Note: IBM highly recommends that a redundant SAN be used with the SAN Volume Controller, however a nonredundant SAN is supported.

To install a SAN Volume Controller into an existing SAN that will be in use during installation, you must first ensure that the switch zoning is set to isolate the new SAN Volume Controller connections from the active part of the SAN.

See the following Web site for specific firmware levels and the latest supported hardware:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

- Consider the design of the SAN according to your requirement for high availability.
- Identify the operating system for each host system that will be connected to the SAN Volume Controller, ensuring compatibility and suitability by performing the following steps:
 1. Specify the host bus adapters (HBAs) for each host.
 2. Define the performance requirements.
 3. Determine the total storage capacity.
 4. Determine the storage capacity per host.
 5. Determine the host LUN sizes.
 6. Determine the total number of ports and the bandwidth needed between the host to the SAN Volume Controller.
 7. Determine if your SAN has enough ports to connect all hosts and backend storage.
 8. Determine if your SAN provides enough ports to connect your backend storage.
- Ensure that the existing SAN components meet the requirements for the SAN Volume Controller by performing the following steps:
 1. Determine the host system versions.
 2. Ensure that the HBAs, switches, and controllers are at or above the minimum requirements.
 3. Identify any components that must be upgraded.

Switch zoning for the SAN Volume Controller

Consider these constraints when you zone a switch.

Overview

The number of virtual paths to each virtual disk is limited. Implementation of the following rules can help you achieve the correct number of virtual paths.

- Each host (or partition of a host) can have between one and four fibre-channel ports.
- Brocade and McData switches can be configured in Vendor Interoperability Mode or in Native Mode. Cisco switches are not currently supported in Vendor Interoperability Mode.
- The default timeout values for fabric is the following:
 - E_A_TOV=10
 - E_D_TOV=2

Operation with values other than these specified is unsupported. IBM recommends that you manually set the domain IDs prior to building the multiswitch fabric and prior to zoning for the following reasons:

- When two switches are joined while active, they can determine if the domain ID is already in use as before. If there is a conflict the domain ID cannot be changed in an active switch. This conflict will cause the fabric merging process to fail.
- The domain ID is used to identify switch ports when zoning is implemented using the domain and switch port number. If domain IDs are negotiated at every fabric

start up, there is no guarantee that the same switch will have the same ID the next time. Therefore, zoning definitions can become invalid.

- If the domain ID is changed after a SAN is set up, some host systems may have difficulty logging back in with the switch, and it might be necessary to reconfigure the host in order to detect devices on the switch again.

The maximum number of paths from the SAN Volume Controller nodes to a host is eight. The maximum number of host bus adapter (HBA) ports is four (for example, no more than two two-port HBAs or four one-port HBAs).

Consider the SAN environment in the following example:

- Two SAN Volume Controller nodes, nodes A and B
- Nodes A and B have four ports each
 1. Node A has ports A0, A1, A2, and A3
 2. Node B has ports B0, B1, B2, and B3
- Four hosts called P, Q, R, and S
- Each of the four hosts has four ports, as described in Table 6.

Table 6. Four hosts and their ports

P	Q	R	S
P0	Q0	R0	S0
P1	Q1	R1	S1
P2	Q2	R2	S2
P3	Q3	R3	S3

- Two switches called X and Y
- One storage controller
- The storage controller has four ports on it called I0, I1, I2, and I3

An example configuration would be the following:

1. Attach ports 1 (A0, B0, P0, Q0, R0, and S0) and 2 (A1, B1, P1, Q1, R1, and S1) of each node and host to switch X.
2. Attach ports 3 (A2, B2, P2, Q2, R2, and S2) and 4 (A3, B3, P3, Q3, R3, and S3) of each node and host to switch Y.
3. Attach ports 1 and 2 (I0 and I1) of the storage controller to switch X.
4. Attach ports 3 and 4 (I2 and I3) of the storage controller to switch Y.

On switch X we would create the following host zones:

5. Create a host zone containing ports 1 (A0, B0, P0, Q0, R0, and S0) of each node and host.
6. Create a host zone containing ports 2 (A1, B1, P1, Q1, R1, and S1) of each node and host.

Similarly, on switch Y we would create the following host zones:

7. Create a host zone on switch Y containing ports 3 (A2, B2, P2, Q2, R2, and S2) of each node and host.
8. Create a host zone on switch Y containing ports 4 (A3, B3, P3, Q3, R3, and S3) of each node and host.

Last, we would create the following storage zone:

9. Create a storage zone that is configured on each switch. Each storage zone contains all the SAN Volume Controller and storage ports on that switch.

In the following example, the SAN environment is similar to the first example, with two additional hosts with two ports each.

- Two SAN Volume Controller nodes called A and B
- Nodes A and B have four ports each
 1. Node A has ports A0, A1, A2, and A3
 2. Node B has ports B0, B1, B2, and B3
- Six hosts called P, Q, R, S, T and U
- Four hosts have four ports each, and two hosts have four ports each as described in Table 7.

Table 7. Six hosts and their ports

P	Q	R	S	T	U
P0	Q0	R0	S0	T0	U0
P1	Q1	R1	S1	T1	U1
P2	Q2	R2	S2	—	—
P3	Q3	R3	S3	—	—

- Two switches called X and Y
- One storage controller
- The storage controller has four ports on it called I0, I1, I2, and I3

An example configuration would be the following:

1. Attach ports 1 (A0, B0, P0, Q0, R0, S0 and T0) and 2 (A1, B1, P1, Q1, R1, S1 and T1) of each node and host to switch X.
2. Attach ports 3 (A2, B2, P2, Q2, R2, S2 and T2) and 4 (A3, B3, P3, Q3, R3, S3 and T3) of each node and host to switch Y.
3. Attach ports 1 and 2 (I0 and I1) of the storage controller to switch X.
4. Attach ports 3 and 4 (I2 and I3) of the storage controller to switch Y.

Attention: Hosts T and U (T0 and U0) and (T1 and U1) are zoned to different SAN Volume Controller ports so that each SAN Volume Controller port is zoned to the same number of host ports.

On switch X we would create the following host zones:

5. Create a host zone containing ports 1 (A0, B0, P0, Q0, R0, S0 and T0) of each node and host.
6. Create a host zone containing ports 2 (A1, B1, P1, Q1, R1, S1 and U0) of each node and host.

Similarly, on switch Y we would create the following host zones:

7. Create a host zone on switch Y containing ports 3 (A2, B2, P2, Q2, R2, S2 and T1) of each node and host.
8. Create a host zone on switch Y containing ports 4 (A3, B3, P3, Q3, R3, S3 and U1) of each node and host.

Last, we would create the following storage zone:

9. Create a storage zone configured on each switch. Each storage zone contains all the SAN Volume Controller and storage ports on that switch.

Related reference

“Fibre-channel switches” on page 76

Follow these guidelines for configuring the fibre-channel switches that are supported on the SAN.

Zoning considerations for Metro Mirror

Consider these constraints when you zone a switch to support the Metro Mirror service.

SAN configurations that use the Metro Mirror feature between two clusters require the following additional switch zoning considerations:

- Additional zones for Metro Mirror. For Metro Mirror operations involving two clusters, these clusters must be zoned so that the nodes in each cluster can see the ports of the nodes in the other cluster.
- Use of extended fabric settings in a switched fabric.
- Use of Inter Switch Link (ISL) trunking in a switched fabric.
- Use of redundant fabrics.

Note: These considerations do not apply if the simpler, intracluster mode of Metro Mirror operation is in use, when only a single cluster is needed.

For intracluster Metro Mirror relationships, no additional switch zones are required. For intercluster Metro Mirror relationships, you must perform the following steps:

1. Form a SAN that contains both clusters that are to be used in the Metro Mirror relationships. If cluster A is in SAN A originally, and cluster B is in SAN B originally, there must be at least one fibre-channel connection between SAN A and SAN B. This connection will be one or more inter-switch links. The fibre-channel switch ports that are associated with these inter-switch ports must not appear in any zone.
2. A single SAN can only be formed out of combining SAN A and SAN B if the domain numbers of the switches in each SAN are different, prior to the connection of the two SANs. Ensure that each switch has a different domain ID before you connect the two SANs.
3. Once the switches in SAN A and SAN B are connected, configure them so that they operate as a single group of switches. Each cluster must retain the same set of zones that were required to operate in the original single SAN configuration.
4. Add a new zone that contains all the switch ports that are connected to SAN Volume Controller ports. This zone contains switch ports that were originally in SAN A and in SAN B.
5. You can adjust the switch zoning so that the hosts originally in SAN A can see cluster B. This allows a host to examine data in both the local and remote cluster if required. This view of both clusters is purely optional and in some cases can complicate the way you operate the overall system. Therefore, unless it is specifically needed, this should not be implemented.
6. Verify that the switch zoning is such that cluster A cannot see any of the back-end storage that is owned by cluster B. Two clusters cannot share the same back-end storage devices.

The following zones are needed in a typical intercluster Metro Mirror configuration:

- A zone in the local cluster that contains all the ports in the SAN Volume Controller nodes in that local cluster and the ports on the backend storage that are associated with that local cluster. These zones are required whether or not Metro Mirror is in use.

- A zone in the remote cluster that contains all the ports in the SAN Volume Controller nodes in that remote cluster and the ports on the back-end storage that are associated with that remote cluster. These zones are required whether or not Metro Mirror is in use.
- A zone that contains all the ports in the SAN Volume Controller nodes in both the local and remote cluster. This zone is required for intercluster communication and is specifically required by Metro Mirror operations.
- Additional zones that contain ports in host HBAs and selected ports on the SAN Volume Controller nodes in a particular cluster. These are the zones that allow a host to see VDisks presented by an I/O group in a particular cluster. These zones are required whether or not Metro Mirror were in use.

Note:

1. While it is normal to zone a server connection so that it is only visible to the local or remote cluster, it is also possible to zone the server so that the host HBA can see nodes in both the local and remote cluster at the same time.
2. Intracluster Metro Mirror operation does not require any additional zones, over and above those needed to run the cluster itself.

Switch operations over long distances

Some SAN switch products provide features that allow the users to tune the performance of I/O traffic in the fabric in a way that can affect Metro Mirror performance.

The two most significant features are ISL trunking and extended fabric.

ISL trunking	<p>Trunking enables the switch to use two links in parallel and still maintain frame ordering. It does this by routing all traffic for a given destination over the same route even when there may be more than one route available. Often trunking is limited to certain ports or port groups within a switch. For example, in the IBM 2109-F16 switch, trunking can only be enabled between ports in the same quad (for example, same group of four ports). For more information on trunking with the MDS, refer to "Configuring Trunking" on the Cisco Systems Web site.</p> <p>Some switch types can impose limitations on concurrent use of trunking and extended fabric operation. For example, with the IBM 2109-F16 switch, it is not possible to enable extended fabric for two ports in the same quad. Thus, extended fabric and trunking are effectively mutually exclusive. (Although it is possible to enable extended fabric operation one link of a trunked pair, this does not offer any performance advantages and adds complexity to the configuration setup. This mixed mode of operation is therefore not recommended.)</p>
---------------------	---

Extended fabric	<p>Extended fabric operation allocates extra buffer credits to a port. This is important over long links usually found in intercluster Metro Mirror operation. Because the time it takes for a frame to traverse the link, it is possible to have more frames in transmission at any instant in time than would be possible over a short link. The additional buffering is required to allow for the extra frames.</p> <p>For example, the default license for the IBM 2109-F16 switch has two extended fabric options: Normal and Extended Normal.</p> <ul style="list-style-type: none"> • Normal is suitable for short links and Extended Normal is suitable for links up to 10 km long. (With the additional Extended fabric license the user gets two extra options: Medium, up to 10-50 km and Long, 50-100 km.) • The Extended Normal setting gives significantly better performance for the links up to 10 km long. Medium and Long settings are not recommended for use in the intercluster Metro Mirror links that are currently supported.
------------------------	---

Operation of a SAN Volume Controller cluster in SAN fabrics with long distance fibre links

A SAN Volume Controller cluster using SAN fabric switches can connect to application hosts, storage controllers, or other SAN Volume Controller clusters, through the use of short or long wave optical fibre channel connections.

The maximum distance between the cluster and host or the cluster and the storage controller is 300 m for short wave and 10 km for long wave optical connections. Longer distances are supported between clusters that use intercluster Metro Mirror.

Follow these guidelines when you use long wave optical fibre channel connections:

- For disaster recovery, each cluster must be regarded as a single entity. This includes the back-end storage that provides the quorum disks for the cluster. Therefore, the cluster and quorum disks must be co-located. Placing components of a single cluster in different physical locations is not recommended.
- All nodes within a cluster must be located in the same set of racks. There can be a large optical distance between the nodes in the same cluster; however, the nodes must be physically co-located to permit effective service and maintenance.
- All nodes in a cluster must be on the same IP subnet. This allows the nodes to assume the same cluster or service IP address.
- A node must be on the same rack as the uninterruptible power supply from which it receives power.

Note: Splitting cluster operation over a long optical distance is not recommended, as only asymmetric disaster recovery can be provided with substantially reduced performance. Instead, use two cluster configurations for all production disaster recovery systems.

Related reference

“Switch operations over long distances” on page 34

Some SAN switch products provide features that allow the users to tune the performance of I/O traffic in the fabric in a way that can affect Metro Mirror performance.

“Zoning considerations for Metro Mirror” on page 33

Consider these constraints when you zone a switch to support the Metro Mirror service.

Performance of fibre-channel extenders

When you are planning to use fibre-channel extenders, it is important to be aware that the performance of the link to the remote location decreases as the distance to the remote location increases.

For fibre-channel IP extenders, throughput is limited by latency and bit error rates. Typical I/O latency can be expected to be 10 milliseconds per kilometer. Bit error rates vary depending on the quality of the circuit provided.

You must review the total throughput rates that might be expected for your planned configuration with the vendor of your fibre-channel extender and your network provider.

Related reference

“Supported fibre-channel extenders” on page 84

The SAN Volume Controller supports the CNT UltraNet Edge Storage Router, the Cisco MDS 9506 and 9509 Director, and Cisco 9516 Fabric Switches to support synchronous copy services.

Nodes

A SAN Volume Controller node is a single processing unit within a SAN Volume Controller cluster.

For redundancy, nodes are deployed in pairs to make up a cluster. A cluster can have one to four pairs of nodes in it. Each pair of nodes is known as an I/O group. Each node can be in *only* one I/O group. A maximum of four I/O groups each containing two nodes is supported.

At any one time, a single node in the cluster manages configuration activity. This configuration node manages a cache of the configuration information that describes the cluster configuration and provides a focal point for configuration commands. If the configuration node fails, another node in the cluster takes over its responsibilities.

Table 8 describes the operational states of a node.

Table 8. Node state

State	Description
Adding	The node was added to the cluster but is not yet synchronized with the cluster state (see Note).
Deleting	The node is in the process of being deleted from the cluster.
Online	The node is operational, assigned to a cluster, and has access to the fibre-channel SAN fabric.
Offline	The node is not operational. The node was assigned to a cluster but is not available on the fibre-channel SAN fabric. Run the Directed Maintenance Procedures to determine the problem.
Pending	The node is transitioning between states and, in a few seconds, will move to one of the other states.

Table 8. Node state (continued)

State	Description
Note: It is possible that a node can stay in the Adding state for a long time. If this is the case, delete the node and then re-add it. However, you should wait for at least 30 minutes before doing this. If the node that has been added is at a lower code level than the rest of the cluster, the node will be upgraded to the cluster code level, which can take up to 20 minutes. During this time the node will be shown as adding.	

Clusters

All configuration and service is performed at the cluster level.

A cluster can consist of two nodes, with a maximum of eight nodes. Therefore, you can assign up to eight SAN Volume Controller nodes to one cluster.

Some service actions can be performed at node level, but all configuration is replicated across all nodes in the cluster. Because configuration is performed at the cluster level, an IP address is assigned to the cluster instead of each node.

All your configuration and service actions are performed at the cluster level. Therefore, after configuring your cluster, you can take advantage of the virtualization and the advanced features of the SAN Volume Controller.

Cluster state and the configuration node

The cluster state holds all configuration and internal cluster data for the cluster. This cluster state information is held in nonvolatile memory. If the mainline power fails, the two uninterruptible power supplies maintain the internal power long enough for the cluster state information to be stored on the internal disk drive of each node. The read and write cache information is also held in nonvolatile memory. Similarly, if the power fails to a node, configuration and cache data for that node is lost and the partner node attempts to flush the cache. The cluster state is still maintained by the other nodes on the cluster.

Figure 8 on page 38 shows an example cluster containing four nodes. The cluster state shown in the grey box does not actually exist, instead each node holds a copy of the entire cluster state.

The cluster contains a single node that is elected as the configuration node. The configuration node can be thought of as the node that controls the updating of cluster state. For example, a user request is made (item 1), that results in a change being made to the configuration. The configuration node controls updates to the cluster (item 2). The configuration node then forwards the change to all nodes (including Node 1), and they all make the state-change at the same point in time (item 3). Using this state-driven model of clustering ensures that all nodes in the cluster know the exact cluster state at any one time.

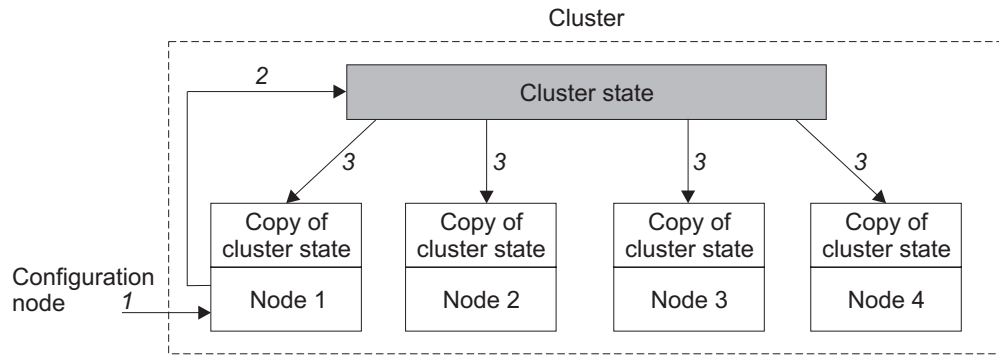


Figure 8. Cluster, nodes, and cluster state.

Cluster state

The cluster state holds all configuration and internal cluster data for the cluster.

The cluster state information is held in non-volatile memory. If the mainline power fails, the two uninterruptible power supply units maintain the internal power long enough for the cluster state information to be stored on the internal SCSI disk drive of each node. The read and write cache information, which is also held in memory, is stored on the internal SCSI disk drives of the nodes in the I/O group that are using that information.

Each node in the cluster maintains an identical copy of the cluster state. When a change is made to the configuration or internal cluster data, then the same change is applied to all nodes. For example, a user configuration request is made to the configuration node. This node forwards the request to all nodes in the cluster and they all make the change to the cluster state at the same point in time. This ensures that all nodes are aware of the configuration change. If the configuration node fails, then the cluster can elect a new node to take over its responsibilities.

Cluster operation and quorum disks

The cluster must contain at least half of its nodes to function.

Nodes are deployed in pairs known as I/O groups, and one to four I/O groups comprise a cluster. In order to function, one node in each I/O group must be operational. If both of the nodes in an I/O group are not operational, access is lost to the VDisks that are managed by the I/O group.

Note that the cluster can survive more than half the nodes failing as long as the cluster has stabilized between failures.

A tie-break situation can occur if exactly half the nodes in a cluster fail at the same time, or if the cluster is divided so that exactly half the nodes in the cluster cannot communicate with the other half. For example, in a cluster of four nodes, if any two nodes fail at the same time or any two cannot communicate with the other two, a tie-break exists and must be resolved.

The cluster automatically chooses three managed disks to be *quorum disks* and assigns them quorum indexes of 0, 1, and 2. One of these disks is used to settle a tie-break condition.

If a tie-break occurs, the first half of the cluster to access the quorum disk after the split has occurred locks the disk and continues to operate. The other side stops. This action prevents both sides from becoming inconsistent with each other.

You can change the assignment of quorum disks at any time by issuing the following command:

```
svctask setquorum
```

I/O groups and uninterruptible power supply

Each pair of nodes is known as an *I/O Group*.

Each node can only be in one I/O group. The I/O groups are connected to the SAN so that all backend storage and all application servers are visible to all of the I/O groups. Each pair has the responsibility to serve I/O on a particular virtual disk.

Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller nodes. Virtual disks are also associated with an I/O group. The SAN Volume Controller does not contain any internal battery backup units and therefore must be connected to an uninterruptible power supply to provide data integrity in the event of a cluster wide power failure.

When an application server performs I/O to a virtual disk, it can access the virtual disk with either of the nodes in the I/O group. A virtual disk can have a preferred node specified when it is created. This is specified after the virtual disk is created. This is the node through which a virtual disk can normally be accessed. As each I/O group only has two nodes, the distributed cache in the SAN Volume Controller need only be two-way. When I/O is performed to a virtual disk, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group.

I/O traffic for a particular virtual disk is, at any one time, managed exclusively by the nodes in a single I/O group. Thus, although a cluster can have many nodes within it, the nodes handle I/O in independent pairs. This means that the I/O capability of the SAN Volume Controller scales well, since additional throughput can be obtained by adding additional I/O groups.

Figure 9 on page 40 shows an example I/O group. A write operation from a host is shown (item 1), that is targeted for virtual disk A. This write is targeted at the preferred node, Node 1 (item 2). The write is cached and a copy of the data is made in the partner node, Node 2's cache (item 3). The write is now complete so far as the host is concerned. At some time later the data is written, or destaged, to storage (item 4). The figure also shows two uninterruptible power supply units (1 and 2) correctly configured so that each node is in a different power domain.




Figure 9. I/O group and uninterruptible power supply

When a node fails within an I/O group, the other node in the I/O group takes over the I/O responsibilities of the failed node. Data loss during a node failure is prevented by mirroring the I/O read/write data cache between the two nodes in an I/O group.

If only one node is assigned to an I/O group or if a node has failed in an I/O group, the cache goes into write-through mode. Therefore, any writes for the virtual disks that are assigned to this I/O group are not cached, but are sent directly to the storage device. If both nodes in an I/O group go offline, the virtual disks that are assigned to the I/O group cannot be accessed.

When a virtual disk is created, you must specify the I/O group that will provide access to the virtual disk. However, virtual disks can be created and added to I/O groups that contain offline nodes. I/O access will not be possible until at least one of the nodes in the I/O group is online.

The cluster also provides a *recovery I/O group*, which is used when both nodes in the I/O group have multiple failures. This allows you to move the virtual disks to the recovery I/O group and then into a working I/O group. I/O access is not possible when virtual disks are assigned to the recovery I/O group.

Uninterruptible power supply and power domains

An uninterruptible power supply protects the cluster against power failures.

If the mainline power fails to one or more nodes in the cluster, the uninterruptible power supply maintains the internal power long enough for the cluster state information to be stored on the internal SCSI disk drive of each node.

It is required that each node in the cluster be connected to an uninterruptible power supply. This allows the cluster to continue to work in degraded mode if one uninterruptible power supply fails.

It is very important that the two nodes in an I/O group are not both connected to the same power domain. Each node of an I/O group must be connected to a different uninterruptible power supply. This configuration ensures that the cache and cluster state information is protected against the failure of the uninterruptible power supply or of the mainline power source. If possible, each uninterruptible power supply should be connected to a different power source. Otherwise, a power source failure will result in the I/O group being taken offline.

The following uninterruptible power supply units can be used with the SAN Volume Controller:

- UPS 5125. This unit can support two nodes.
- UPS 5115. This unit can support one node.

These units can be used in combination in a cluster. Table 9 shows the required number of uninterruptible power supply units for the number of nodes in a cluster.

Table 9. Required uninterruptible power supply (UPS) units

Number of nodes	Number of required UPS 5125 units	Number of required UPS 5115 units
2 nodes	2 UPS units	2 UPS units
4 nodes	2 UPS units	4 UPS units
6 nodes	4 UPS units	6 UPS units
8 nodes	4 UPS units	8 UPS units

When nodes are added to the cluster, you must specify the I/O group that they will join. The configuration interfaces will also check the uninterruptible power supply units and ensure that the two nodes in the I/O group are not connected to the same uninterruptible power supply units.

Figure 10 shows a cluster of four nodes, with two I/O groups and two uninterruptible power supply units.




Figure 10. Relationship between I/O groups and uninterruptible power supply units

Attention: Do not connect two clusters to the same pair of uninterruptible power supply units. Both clusters will be lost in the event that a power failure occurs on both of the units.

Disk controllers

A disk controller is 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).

When configuring disk controllers, ensure that you configure and manage the disk controllers and its devices for optimal performance.

The supported RAID controllers are detected by the cluster and reported by the user interfaces. The cluster can determine which MDisks each controller has and can also provide a view of MDisks filtered by controller. This view enables you to associate the MDisks with the RAID arrays that the controller presents.

Note: The SAN Volume Controller supports RAID controllers, but it is possible to configure a controller as a non-RAID controller. RAID controllers provide redundancy at the disk level. Therefore, a single physical disk failure does not cause an MDisk failure, an MDisk group failure, or a failure in the virtual disks (VDisks) that were created from the MDisk group.

The controller can have a local name for the RAID arrays or single disks that it is providing. However it is not possible for the nodes in the cluster to determine this name because the namespace is local to the controller. The controller will surface these disks with a unique ID, the controller LUN or LU number. This ID, along with the controller serial number or numbers (there may be more than one controller), can be used to associate the managed disks in the cluster with the RAID arrays presented by the controllers.

To prevent loss of data, virtualize only those RAID arrays that provide some form of redundancy, that is, RAID 1, RAID 10, RAID 0+1 or RAID 5. Do not use RAID 0 because a single physical disk failure might cause the failure of many VDIs.

Unsupported disk controller systems (generic controllers)

When a disk controller system is detected on the SAN, the SAN Volume Controller attempts to recognize it using its Inquiry data. If the disk controller system is recognized as one of the explicitly supported storage models, then the SAN Volume Controller uses error recovery programs that can be tailored to the known needs of the disk controller system. If the storage controller is not recognized, then the SAN Volume Controller configures the disk controller system as a generic controller. A generic controller might not function correctly when it is addressed by a SAN Volume Controller. The SAN Volume Controller does not regard accessing a generic controller as an error condition and, consequently, does not log an error. MDisks that are presented by generic controllers are not eligible to be used as quorum disks.

Related concepts

“Managed disks” on page 54

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

Data migration

Data migration affects the mapping of the extents for a virtual disk (VDisk) to the extents for a managed disk (MDisk).

The host can access the VDisk during the data migration process.

Applications for data migration

There are several applications for data migration.

- Redistribution of workload within a cluster across managed disks:
 - Moving workload onto newly installed storage
 - Moving workload from old or failing storage, prior to replacing it
 - Moving workload to rebalance workload that has changed
- Migrating data from legacy disks to disks that are managed by the SAN Volume Controller.

Image mode virtual disk migration

Image mode virtual disks (VDisks) have a special property that allows the last extent in the VDisk to be a partial extent.

You can migrate striped and sequential Vdisks to an image mode VDisk.

Copy Services

The SAN Volume Controller provides Copy Services that enable you to copy virtual disks (VDisks).

The following Copy Services options are available for all supported hosts that are connected to the SAN Volume Controller.

FlashCopy

Makes an instant, point-in-time copy from a source VDisk to a target VDisk.

Metro Mirror

Provides a consistent copy of a source VDisk on a target VDisk. Data is written to the target VDisk synchronously after it is written to the source VDisk, so the copy is continuously updated.

Applications for FlashCopy

You can use FlashCopy to back up data that changes frequently. After you create the point-in-time copy, you can back it up to tertiary storage such as tape.

Another use of FlashCopy is for application testing. It is often important and useful to test a new version of an application on real business data before you move the application into production. This reduces the risk that the new application will fail because it is not compatible with actual business data.

You can also use FlashCopy to create copies for your auditing and data mining purposes.

In the scientific and technical arena, FlashCopy can create restart points for long-running batch jobs. Therefore, if a batch job fails many days into its run, you

might be able to restart the job from a saved copy of its data. This is preferable to rerunning the entire multiday job.

Applications for Metro Mirror

Disaster recovery is the primary application for Metro Mirror. Because an exact copy of your business data can be maintained at a remote location, you can use your remote location as a recovery site in the event of a local disaster.

FlashCopy

FlashCopy is a copy service that is available with the SAN Volume Controller.

It copies the contents of a source virtual disk (VDisk) to a target VDisk. Any data that existed on the target disk is lost and is replaced by the copied data. After the copy operation has been completed, the target virtual disks contain the contents of the source virtual disks as they existed at a single point in time unless target writes have been performed. Although the copy operation takes some time to complete, the resulting data on the target is presented in such a way that the copy appears to have occurred immediately. FlashCopy is sometimes described as an instance of a time-zero copy (T 0) or point-in-time copy technology. Although the FlashCopy operation takes some time, this time is several orders of magnitude less than the time that would be required to copy the data using conventional techniques.

It is difficult to make a consistent copy of a data set that is being constantly updated. Point-in-time copy techniques are used to help solve this problem. If a copy of a data set is created using a technology that does not provide point-in-time techniques and the data set changes during the copy operation, then the resulting copy may contain data which is not consistent. For example, if a reference to an object is copied earlier than the object itself and the object is moved before it is itself copied, the copy will contain the referenced object at its new location but the reference will point to the old location.

Source VDIsks and target VDIsks must meet the following requirements:

- They must be the same size.
- The same cluster must manage them.

Related concepts

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

FlashCopy mappings

A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

Because FlashCopy copies one VDisk to another VDisk, the SAN Volume Controller Console must be aware of that relationship. A particular VDisk can take part in only one mapping; that is, a VDisk can be the source or target of only one mapping. You cannot, for example, make the target of one mapping the source of another mapping.

FlashCopy makes an instant copy of a VDisk at the time that it is started. To create a FlashCopy of a virtual disk, you must first create a mapping between the source VDisk (the disk that is copied) and the target VDisk (the disk that receives the copy). The source and target must be of equal size.

To copy a VDisk, it must be part of a FlashCopy mapping or of a consistency group.

A FlashCopy mapping can be created between any two VDisks in a cluster. It is not necessary for the VDisks to be in the same I/O group or managed disk group. When a FlashCopy operation is started, a checkpoint is made of the source VDisk. No data is actually copied at the time a start occurs. Instead, the checkpoint creates a bitmap that indicates that no part of the source VDisk has been copied. Each bit in the bitmap represents one region of the source VDisk. Such a region is called a *grain*.

After a FlashCopy operation starts, read operations to the source VDisk continue to occur. If new data is written to the source (or target) VDisk, then the existing data on the source is copied to the target VDisk before the new data is written to the source (or target) VDisk. The bitmap is updated to mark that the grain of the source VDisk has been copied so that later write operations to the same grain do not recopy the data.

Similarly, during a read operation to the target VDisk the bitmap is used to determine the grain has been copied. If the grain has been copied, the data is read from the target VDisk. If the grain has not been copied, the data is read from the source VDisk.

When you create a mapping, you specify the background copy rate. This rate determines the priority that is given to the background copy process. If you want to end with a copy of the whole source at the target (so that the mapping can be deleted, but the copy can still be accessed at the target), you must copy all the data that is on the source VDisk to the target VDisk.

When a mapping is started and the background copy rate is greater than zero (or a value other than NOCOPY is selected in the SAN Volume Controller Console's Creating FlashCopy Mappings panel), the unchanged data is copied to the target, and the bitmap is updated to show that the copy has occurred. After a time, the length of which depends on the priority given and the size of the VDisk, the whole VDisk is copied to the target. The mapping returns to the idle/copied state. You can restart the mapping at any time to create a new copy at the target; the process copy starts again.

If the background copy rate is zero (or NOCOPY), only the data that changes on the source is copied to the target. The target never contains a copy of the whole source unless every extent is overwritten at the source. You can use this copy rate when you need only a temporary copy of the source.

You can stop the mapping at any time after it has been started. This action makes the target inconsistent and therefore the target VDisk is taken offline. You must restart the mapping to correct the target.

FlashCopy mapping states

At any point in time, a FlashCopy mapping is in one of the following states:

Idle or copied

The source and target VDisks act as independent VDisks even if a FlashCopy mapping exists between the two. Read and write caching is enabled for both the source and the target.

Copying

The copy is in progress.

Prepared

The mapping is ready to start. While in this state, the target VDisk is offline.

Preparing

Any changed write data for the source VDisk is flushed from the cache. Any read or write data for the target VDisk is discarded from the cache.

Stopped

The mapping is stopped because either you issued a command or an input/output (I/O) error occurred. Preparing and starting the mapping again can restart the copy.

Suspended

The mapping started, but it did not complete. The source VDisk might be unavailable, or the copy bitmap might be offline. If the mapping does not return to the copying state, stop the mapping to reset the mapping.

Before you start the mapping, you must prepare it. By preparing the mapping, you ensure that the data in the cache is destaged to disk and that a consistent copy of the source exists on disk. At this time, the cache goes into write-through mode. Data that is written to the source is not cached in the SAN Volume Controllers; it passes straight through to the managed disks. The prepare operation for the mapping might take you a few minutes; the actual length of time depends on the size of the source VDisk. You must coordinate the prepare operation with the operating system. Depending on the type of data that is on the source VDisk, the operating system or application software might also cache data write operations. You must flush, or synchronize, the file system and application program before you prepare for, and finally start, the mapping.

Note: The `svctask startfcmap` command can take some time to process.

For customers who do not need the complexity of consistency groups, the SAN Volume Controller allows a FlashCopy mapping to be treated as an independent entity. In this case the FlashCopy mapping is known as a stand-alone mapping. For FlashCopy mappings which have been configured in this way, the **Prepare** and **Start** commands are directed at the FlashCopy mapping name rather than the consistency group ID.

Veritas Volume Manager

For FlashCopy target VDisks, the SAN Volume Controller sets a bit in the inquiry data for those mapping states where the target VDisk could be an exact image of the source VDisk. Setting this bit enables the Veritas Volume Manager to distinguish between the source and target VDisks and thus provide independent access to both.

Related concepts

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

FlashCopy consistency groups

A consistency group is a container for mappings. You can add many mappings to a consistency group.

The consistency group is specified when the mapping is created. You can also change the consistency group later. When you use a consistency group, you

prepare and trigger that group instead of the various mappings. This ensures that a consistent copy is made of all the source virtual disks (VDisks). Do not place mappings that you want to control at an individual level instead of at a consistency group level into a consistency group. These individual mappings are known as stand-alone mappings.

To copy a VDisk, it must be part of a FlashCopy mapping or of a consistency group.

When you copy data from one VDisk to another, that data might not include all that you need to enable you to use the copy. Many applications have data that spans multiple VDIsks and that include the requirement that data integrity is preserved across VDIsks. For example, the logs for a particular database usually reside on a different VDisk than the VDisk that contains the data.

Consistency groups address the problem when applications have related data that spans multiple VDIsks. In this situation, FlashCopy must be performed in a way that preserves data integrity across the multiple VDIsks. One requirement for preserving the integrity of data being written is to ensure that dependent writes are run in the intended sequence of the application.

FlashCopy consistency group states

At any point in time, a FlashCopy consistency group is in one of the following states:

Idle or copied

The source and target VDIsks act independently even if a FlashCopy consistency group exists. Read and write caching is enabled for the source VDIsks and target VDIsks.

Copying

The copy is in progress.

Prepared

The consistency group is ready to start. While in this state, the target VDIsks are offline.

Preparing

Any changed write data for the source VDIsks is flushed from the cache.
Any read or write data for the target VDIsks is discarded from the cache.

Stopped

The consistency group is stopped because either you issued a command or an input/output (I/O) error occurred. Preparing and starting the consistency group again can restart the copy.

Suspended

The consistency group was started, but it did not complete. The source VDIsks might be unavailable, or the copy bitmap might be offline. If the consistency group does not return to the copying state, stop the consistency group to reset the consistency group.

Related concepts

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Metro Mirror

Metro Mirror (formerly known as Remote Copy) enables you to set up a relationship between two virtual disks, so that updates that are made by an application to one virtual disk are mirrored on the other virtual disk.

Although the application only writes to a single virtual disk (VDisk), the SAN Volume Controller maintains two copies of the data. If the copies are separated by a significant distance, the Metro Mirror copy can be used as a backup for disaster recovery. A prerequisite for the SAN Volume Controller Metro Mirror operations between two clusters is that the SAN fabric to which they are attached provides adequate bandwidth between the clusters.

One VDisk is designated the primary and the other VDisk is designated the secondary. Host applications write data to the primary VDisk, and updates to the primary VDisk are copied to the secondary VDisk. Normally, host applications do not perform input or output operations to the secondary VDisk. When a host writes to the primary VDisk, it will not receive confirmation of I/O completion until the write operation has completed for the copy on the secondary disk, as well as on the primary.

Metro Mirror supports the following features:

- Intracluster copying of a VDisk, in which both VDIsks belong to the same cluster and I/O group within the cluster.
- Intercluster copying of a VDisk, in which one VDisk belongs to a cluster and the other VDisk belongs to a different cluster

Note: A cluster can only participate in active Metro Mirror relationships with itself and one other cluster.

- Intercluster and intracluster Metro Mirror can be used concurrently within a cluster.
- The intercluster link is bidirectional. This means that it can copy data from clusterA to clusterB for one pair of VDIsks while copying data from clusterB to clusterA for a different pair of VDIsks.
- The copy direction can be reversed for a consistent relationship by issuing a simple **switch** command. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.
- Metro Mirror consistency groups are supported for ease of managing a group of relationships that need to be kept in sync for the same application. This also simplifies administration, as a single command issued to the consistency group will be applied to all the relationships in that group.

Related concepts

“IBM TotalStorage Metro Mirror for SAN Volume Controller” on page 49
IBM TotalStorage Metro Mirror for SAN Volume Controller (formerly known as synchronous Remote Copy) provides a *consistent* copy, which means that the primary virtual disk (VDisk) is always the exact match of the secondary VDisk.

“Metro Mirror consistency groups” on page 49

Metro Mirror provides the facility to group a number of relationships into a Metro Mirror consistency group so that they can be manipulated in unison.

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

IBM TotalStorage Metro Mirror for SAN Volume Controller

IBM TotalStorage Metro Mirror for SAN Volume Controller (formerly known as synchronous Remote Copy) provides a *consistent copy*, which means that the primary virtual disk (VDisk) is always the exact match of the secondary VDisk.

The host application writes data to the primary VDisk but does not receive the final status on the write operation until the data is written to the secondary VDisk. For disaster recovery, this mode is the only practical mode of operation because a consistent copy of the data is maintained. However, Metro Mirror is slower than Global Mirror because of the latency time and bandwidth limitations that are imposed by the communication link to the secondary site.

Related concepts

“Metro Mirror consistency groups”

Metro Mirror provides the facility to group a number of relationships into a Metro Mirror consistency group so that they can be manipulated in unison.

“Virtual disks” on page 59

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Metro Mirror consistency groups

Metro Mirror provides the facility to group a number of relationships into a Metro Mirror consistency group so that they can be manipulated in unison.

Certain uses of Metro Mirror require the manipulation of more than one relationship. A command issued to the consistency group will be applied to all of the relationships in the group simultaneously.

For some uses it might be that the relationships share some loose association and that the grouping simply provides a convenience for the administrator. But a more significant use arises when the relationships contain virtual disks (VDisks) that have a tighter association. One example is when the data for an application is spread across more than one VDisk. A more complex example is when multiple applications run on different host systems. Each application has data on different VDIs, and these applications exchange data with each other. Both these examples are cases in which specific rules exist as to how the relationships must be manipulated, in unison. This ensures that the set of secondary VDIs contains usable data. The key property is that these relationships be consistent. Hence, the groups are called consistency groups.

A relationship can be part of a single consistency group or not be part of a consistency group at all. Relationships that are not part of a consistency group are called stand-alone relationships. A consistency group can contain zero or more relationships. All the relationships in a consistency group must have matching master and auxiliary clusters. All relationships in a consistency group must also have the same copy direction and state.

Metro Mirror consistency group states

Inconsistent (Stopped)

The primary VDIs are accessible for read and write input/output (I/O) operations but the secondary VDIs are not accessible for either. A copy process must be started to make the Secondary VDIs consistent.

Inconsistent (Copying)

The primary VDIs are accessible for read and write I/O operations but the

| secondary VDisk are not accessible for either. This state is entered after a
| **Start** command is issued to a consistency group in the InconsistentStopped
| state. This state is also entered when a **Start** command is issued, with the
| force option, to a consistency group in the Idling or ConsistentStopped
| state.

| **Consistent (Stopped)**

| The secondary VDIsks contain a consistent image, but it might be
| out-of-date with respect to the primary VDIsks. This state can happen when
| a relationship was in the ConsistentSynchronized state and experiences an
| error which forces a freeze of the consistency group. This state can also
| happen when a relationship is created with the CreateConsistentFlag set to
| TRUE.

| **Consistent (Synchronized)**

| The primary VDIsks are accessible for read and write I/O operations. The
| secondary VDIsks are accessible for read-only I/O operations.

| **Idling** Master VDIsks and Auxiliary VDIsks are operating in the primary role.
| Consequently the VDIsks are accessible for write I/O operations.

| **Idling (Disconnected)**

| The VDIsks in this half of the consistency group are all operating in the
| primary role and can accept read or write I/O operations.

| **Inconsistent (Disconnected)**

| The VDIsks in this half of the consistency group are all operating in the
| secondary role and will not accept read or write I/O operations.

| **Consistent (Disconnected)**

| The VDIsks in this half of the consistency group are all operating in the
| secondary role and will accept read I/O operations but not write I/O
| operations

| **Empty**

| The consistency group contains no relationships.

| **Related concepts**

| “Virtual disks” on page 59

| A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage
| area network (SAN).

Chapter 5. Object descriptions

The SAN Volume Controller is based on a group of virtualization concepts. Before setting up the system, you should understand the concepts and the objects in the system.

The smallest processing unit in a SAN Volume Controller is a single **node**. Nodes are deployed in pairs to make up a cluster. A cluster can consist of one to four pairs of nodes. Each pair of nodes is known as an **I/O group**. Each node may be in only one I/O group.

Virtual disks (Vdisks) are logical disks that are presented by clusters. Each virtual disk is associated with a particular I/O group. The nodes in the I/O group provide access to the virtual disks in the I/O group. When an application server performs I/O to a virtual disk, it can access the virtual disk with either of the nodes in the I/O group. Because each I/O group only has two nodes, the distributed cache that the SAN Volume Controller provides is only two-way.

Each node does not contain any internal battery backup units and therefore must be connected to an **uninterruptible power supply (UPS)** which provides data integrity in the event of a cluster wide power failure. In such situations, the UPS maintains power to the nodes while the contents of the distributed cache are dumped to an internal drive.

The nodes in a cluster see the storage presented by backend **disk controllers** as a number of disks, known as **managed disks (MDisks)**. Because the SAN Volume Controller does not attempt to provide recovery from physical disk failures within the backend disk controllers, a managed disk is usually, but not necessarily, a RAID array.

Each managed disk is divided up into a number of **extents** which are numbered, from 0, sequentially from the start to the end of the managed disk.

Managed disks are collected into groups, known as **managed disk groups (MDisk group)**. Virtual disks are created from the extents contained by a managed disk group. The managed disks that constitute a particular virtual disk must all come from the same managed disk group.

At any one time, a single node in the cluster can manage configuration activity. This **configuration node** manages a cache of the information that describes the cluster configuration and provides a focal point for configuration.

The SAN Volume Controller detects the fibre-channel ports that are connected to the SAN. These correspond to the Worldwide Port Names (WWPNs) of the host bus adapter (HBA) fibre-channels that are present in the application servers. The SAN Volume Controller allows you to create logical host objects that group WWPNs that belong to a single application server or to a set of them.

Application servers can only access virtual disks that have been allocated to them. Virtual disks can be mapped to a host object. Mapping a virtual disk to a host object makes the virtual disk accessible to the WWPNs in that host object, and hence the application server itself.

The SAN Volume Controller provides block-level aggregation and volume management for disk storage within the SAN. In simpler terms, this means that the

SAN Volume Controller manages a number of back-end storage controllers and maps the physical storage within those controllers into logical disk images that can be seen by application servers and workstations in the SAN. The SAN is configured in such a way that the application servers cannot see the back-end physical storage. This prevents any possible conflict between the SAN Volume Controller and the application servers both trying to manage the back-end storage.

Figure 11 illustrates the objects that are described in this section and their logical positioning in a virtualized system. To simplify the example, virtual disk to host mappings are not shown.

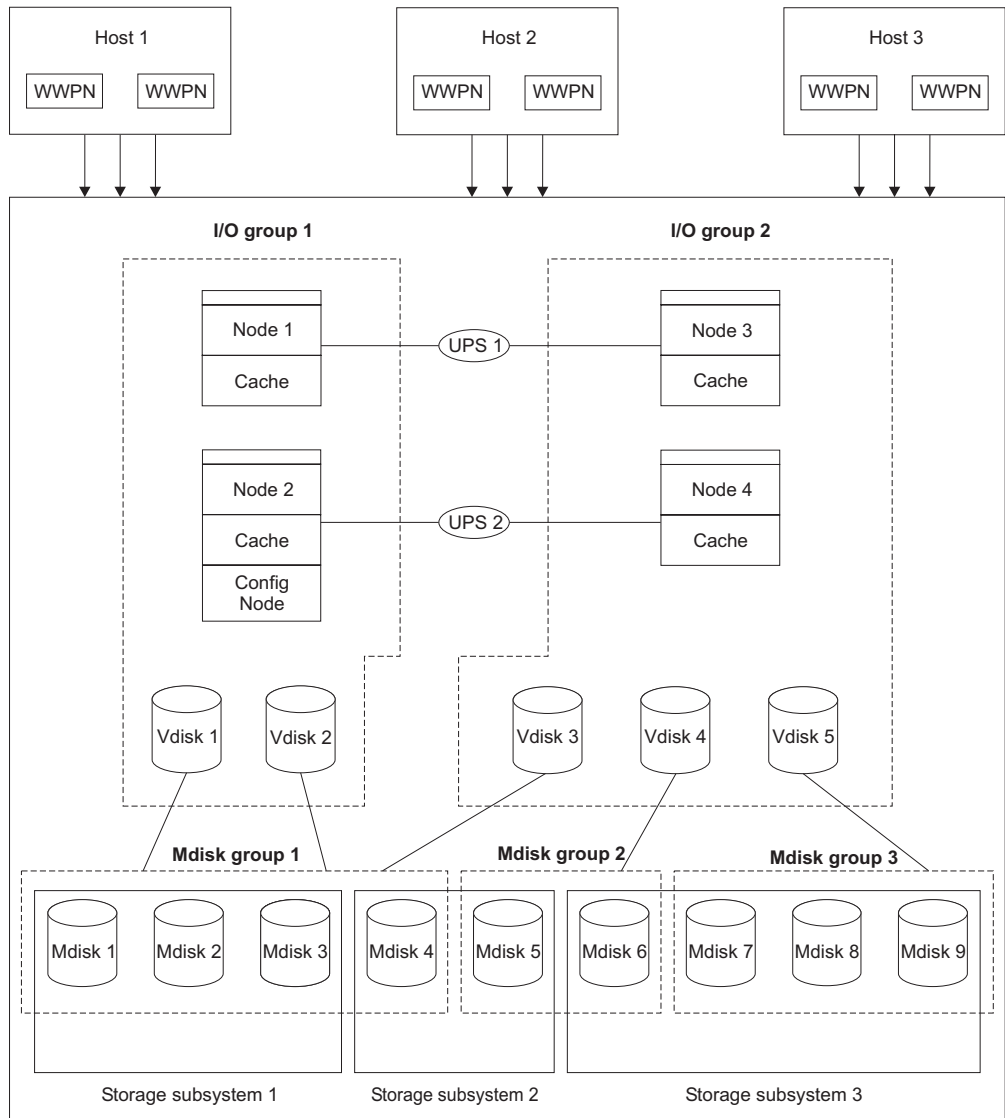


Figure 11. Objects in a virtualized system

Storage subsystems

A storage subsystem is 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.

Storage subsystems that are attached to the SAN fabric provide the physical storage devices that the cluster detects as managed disks. These are usually RAID arrays, because the SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. The nodes in the cluster are connected to one or more fibre-channel SAN fabrics.

The exported storage devices are detected by the cluster and reported by the user interfaces. The cluster can also determine which managed disks each storage subsystem is presenting, and can provide a view of managed disks filtered by the storage subsystem. This allows you to associate the managed disks with the RAID arrays that the subsystem exports.

The storage subsystem can have a local name for the RAID arrays or single disks that it is providing. However it is not possible for the nodes in the cluster to determine this name, because the namespace is local to the storage subsystem. The storage subsystem will surface these storage devices with a unique ID, the logical unit number (LUN). This ID, along with the storage subsystem serial number or numbers (there may be more than one controller in a storage subsystem), can be used to associate the managed disks in the cluster with the RAID arrays exported by the subsystem.

Storage subsystems export storage to other devices on the SAN. The physical storage associated with a subsystem is normally configured into RAID arrays that provide recovery from physical disk failures. Some subsystems also allow physical storage to be configured as RAID-0 arrays (striping) or as JBODs. However, this does not provide protection against a physical disk failure and, with virtualization, can lead to the failure of many virtual disks.

Many storage subsystems allow the storage provided by a RAID array to be divided up into many SCSI logical units (LUs) that are presented on the SAN. With the SAN Volume Controller, it is recommended that storage subsystems are configured to present each RAID array as a single SCSI LU that will be recognized by the SAN Volume Controller as a single managed disk. The virtualization features of the SAN Volume Controller can then be used to divide up the storage into virtual disks.

Some storage subsystems allow the exported storage to be increased in size. The SAN Volume Controller does not use this extra capacity. Instead of increasing the size of an existing managed disk, add a new managed disk to the managed disk group and the extra capacity will be available for the SAN Volume Controller to use.

Attention: If you delete a RAID that is being used by the SAN Volume Controller, the MDisk group goes offline and the data in that group is lost.

When you are configuring your storage subsystems, ensure that you configure and manage your subsystems and its devices for optimal performance.

The cluster detects and provides a view of the storage subsystems that the SAN Volume Controller supports. The cluster can also determine which MDisks each subsystem has and can provide a view of MDisks filtered by device. This view enables you to associate the MDisks with the RAID arrays that the subsystem presents.

Note: The SAN Volume Controller Console supports storage that is internally configured as a RAID array. However, it is possible to configure a storage subsystem as a non-RAID device. RAID provides redundancy at the disk level. For RAID devices, a single physical disk failure does not cause an

MDisk failure, an MDisk group failure, or a failure in the virtual disks (VDisks) that were created from the MDisk group.

Storage subsystems reside on the SAN fabric and are addressable by one or more fibre-channel ports (target ports). Each port has a unique name known as a worldwide port name (WWPN).

Managed disks

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

A managed disk might, therefore, consist of multiple physical disks that are presented as a single logical disk to the SAN. A managed disk always provides usable blocks of physical storage to the cluster even if it does not have a one-to-one correspondence with a physical disk.

Each managed disk is divided into a number of *extents*, which are numbered, from 0, sequentially from the start to the end of the managed disk. The extent size is a property of managed disk groups. When an MDisk is added to an MDisk group, the size of the extents that the MDisk will be divided into depends on the attribute of the MDisk group to which it has been added.

Access modes

The access mode determines how the cluster uses the MDisk. These are the possible modes:

Unmanaged

The MDisk is not used by the cluster.


Managed

The MDisk is assigned to an MDisk group and provides extents that virtual disks (VDisks) can use.

Image The MDisk is assigned directly to a VDisk with a one-to-one mapping of extents between the MDisk and the VDisk.

Attention: If you add a managed disk that contains existing data to a managed disk group, you lose the data that it contains. The *image mode* is the only mode that will preserve this data.

Figure 12 on page 55 shows physical disks and managed disks.





Key:  = Physical disks  = Logical disks (managed disks as seen by the 2145)

Figure 12. Controllers and MDisks

Table 10 describes the operational states of a managed disk:

Table 10. Managed disk status

Status	Description
Online	The MDisk can be accessed by all online nodes. That is, all the nodes that are currently working members of the cluster can access this MDisk. The MDisk is online when the following conditions are met: <ul style="list-style-type: none"> • All timeout error recovery procedures complete and report the disk as online. • LUN inventory of the target ports correctly reported the MDisk. • Discovery of this LUN completed successfully. • All of the managed disk target ports report this LUN as available with no fault conditions.
Degraded	The MDisk cannot be accessed by all the online nodes. That is, one or more (but not all) of the nodes that are currently working members of the cluster cannot access this MDisk. The MDisk may be partially excluded; that is, some of the paths to the MDisk (but not all) have been excluded.
Excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem. You can reset an MDisk and include it in the cluster again by running the svctask includemdisk command.
Offline	The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN, the storage subsystem, or one or more physical disks connected to the storage subsystem. The MDisk will only be reported as offline if all paths to the disk fail.

Extents

Each MDisk is divided into chunks of equal size called *extents*. Extents are a unit of mapping the data between MDisks and virtual disks (VDisks).

Attention: If you have observed intermittent breaks in links or if you have been replacing cables or connections in the storage area network fabric, you may have one or more MDisks in degraded status. If an I/O operation is attempted when a link is broken and the I/O operation fails several times, then the system partially excludes the MDisk and it changes the status of the MDisk to degraded. You must include the MDisk to resolve the problem. You can include the MDisk by either selecting the Include MDisk task from the Work with Managed Disks - Managed Disk panel in the SAN Volume Controller Console, or by issuing the following command:

```
svctask includemdisk <mdiskname/id>
```

Managed disk path Each managed disk has an online path count, which is the number of nodes that have access to that managed disk; this represents a summary of the I/O path status between the cluster nodes and the particular storage device. The maximum path count is the maximum number of paths that have been detected by the cluster at any point in the past. Thus, if the current path count is not equal to the maximum path count, then the particular managed disk might be degraded. That is, one or more nodes might not see the managed disk on the fabric.

Managed disk groups

An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

Figure 13 shows an MDisk group containing four MDisks.




Figure 13. MDisk group

All MDisks in a group are split into extents of the same size. VDisks are created from the extents that are available in the group. You can add MDisks to an MDisk group at any time. This way you increase the number of extents that are available for new VDisks or to expand existing VDisks.

Note: RAID array partitions on HP StorageWorks subsystems controllers are only supported in single-port attach mode. MDisk groups that consist of single-port attached subsystems and other storage subsystems are not supported.

You can add MDisks to an MDisk group at any time either to increase the number of extents that are available for new VDIs or to expand existing VDIs. You can add only MDisks that are in unmanaged mode. When MDisks are added to a group, their mode changes from unmanaged to managed.

You can delete MDisks from a group under the following conditions:

- VDIs are not using any of the extents that are on the MDisk.
- Enough free extents are available elsewhere in the group to move any extents that are in use from this MDisk.

Attention: If you delete an MDisk group, you destroy all the VDIs that are made from the extents that are in the group. If the group is deleted, you cannot recover the mapping that existed between extents that are in the group or the extents that the VDIs use. The MDisks that were in the group are returned to unmanaged mode and can be added to other groups. Because the deletion of a group can cause a loss of data, you must force the deletion if VDIs are associated with it.

Table 11 describes the operational states of an MDisk group.

Table 11. Managed disk group status

Status	Description
Online	The MDisk group is online and available. All the MDisks in the group are available.
Degraded	The MDisk group is available; however, one or more nodes cannot access all the MDisks in the group.
Offline	The MDisk group is offline and unavailable. No nodes in the cluster can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.

Attention: If a single MDisk in an MDisk group is offline and therefore cannot be seen by any of the online nodes in the cluster, then the MDisk group of which this MDisk is a member goes offline. This causes *all* the VDIsks that are being presented by this MDisk group to go offline. Take care when you create MDisk groups to ensure an optimal configuration.

Consider the following guidelines when you create MDisk groups:

- If you are creating image-mode VDIsks, do not put all of these VDIsks into one MDisk group because a single MDisk failure results in all of these VDIsks going offline. Allocate your image-mode VDIsks between your MDisk groups.
- Ensure that all MDisks that are allocated to a single MDisk group are the same RAID type. This ensures that a single failure of a physical disk in the storage subsystem does not take the entire group offline. For example, if you have three RAID-5 arrays in one group and add a non-RAID disk to this group, you lose access to all the data striped across the group if the non-RAID disk fails. Similarly, for performance reasons, you must not mix RAID types. The performance of all VDIsks will be reduced to the lowest performer in the group.
- If you intend to keep the virtual disk allocation within the storage exported by a storage subsystem, ensure that the MDisk group that corresponds with a single subsystem is presented by that subsystem. This also enables nondisruptive migration of data from one subsystem to another subsystem and simplifies the decommissioning process if you want to decommission a controller at a later time.
- Except when you migrate between groups, you must associate a VDisk with just one MDisk group.
- An MDisk can be associated with just one MDisk group.

Extents

To track the space that is available on an MDisk, the SAN Volume Controller divides each MDisk into chunks of equal size. These chunks are called *extents* and are indexed internally. Extent sizes can be 16, 32, 64, 128, 256, or 512 MB.

You specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group.

MDisk groups can have different extent sizes, however this is not recommended. Different extent sizes place restrictions on the use of data migration. You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes.

You can use Copy Services to copy a VDisk between MDisk groups that have different extent sizes. You have the following options:

- Use FlashCopy to copy a VDisk between a source and a destination MDisk group that have different extent sizes.
- Use intracluster Metro Mirror to copy a VDisk between a source and a destination MDisk group that have different extent sizes.

The choice of extent size affects the total amount of storage that will be managed by the cluster. Table 12 on page 59 shows the maximum amount of storage that can be managed by a cluster for each extent size.

Table 12. Capacities of the cluster given extent size

Extent size	Maximum storage capacity of cluster
16 MB	64 TB
32 MB	128 TB
64 MB	256 TB
128 MB	512 TB
256 MB	1 PB
512 MB	2 PB

A cluster can manage 4 million extents ($4 \times 1024 \times 1024$). The choice of extent size determines the amount of storage that will be managed by a cluster. For example, with a 16MB extent size, the cluster can manage up to $16\text{MB} \times 4\text{M} = 64\text{TB}$ of storage.

When choosing an extent size, consider your future needs. For example, if you currently have 40TB of storage, and you specify an extent size of 16MB, the capacity of the MDisk group is limited to 64TB of storage in the future. If you select an extent size of 64MB, the capacity of the MDisk group is 256TB.

Using a larger extent size can waste storage. When a VDisk is created, the storage capacity for the VDisk is rounded to a whole number of extents. If you configure the system to have a large number of small VDIsks and you use a large extent size, then this can cause storage to be wasted at the end of each VDisk.

Virtual disks

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Application servers on the SAN access VDIsks, not managed disks (MDisks). VDIsks are created from a set of extents in an MDisk group. There are three types of VDIsks: striped, sequential, and image.

Types

You can create the following types of VDIsks:

Striped

The striping is at extent level. One extent is allocated, in turn, from each managed disk that is in the group. For example, a managed disk group that has 10 MDisks takes one extent from each managed disk. The 11th extent is taken from the first managed disk, and so on. This procedure, known as a round-robin, is similar to RAID-0 striping.

You can also supply a list of MDisks to use as the stripe set. This list can contain two or more MDisks from the managed disk group. The round-robin procedure is used across the specified stripe set.

Attention: Take care when you specify a stripe set if your MDisk group contains MDisks of unequal size. By default, striped VDIs are striped across all MDisks in the group. If some of the MDisks are smaller than others, the extents on the smaller MDisks will be used up before the larger MDisks run out of extents. Manually specifying the stripe set in this case might result in the VDisk not being created.

If you are unsure there is sufficient free space to create a striped VDisk select one of the following options:

- Check the free space on each MDisk in the group, using the **svcinfolsfreeextents** command
- Let the system automatically create the VDisk, by not supplying a specific stripe set.

Figure 14 shows an example of a managed disk group containing three MDisks. This figure also shows a striped virtual disk created from the extents available in the group.

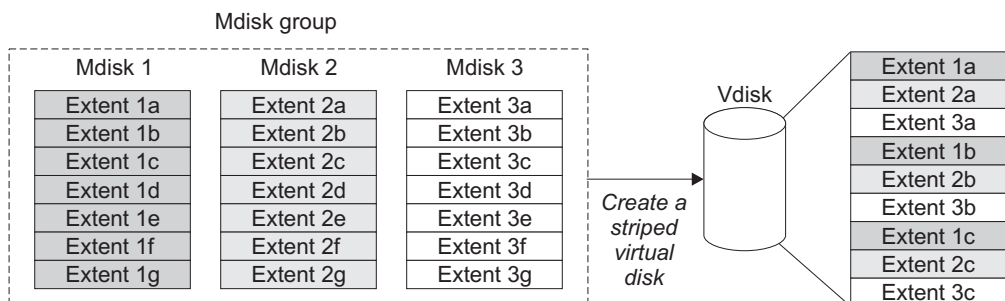


Figure 14. Managed disk groups and VDIs

Sequential

When selected, extents are allocated sequentially on one managed disk to create the virtual disk if enough consecutive free extents are available on the chosen managed disk.

Image Image-mode VDIs are special VDIs that have a direct relationship with one managed disk. If you have a managed disk that contains data that you want to merge into the cluster, you can create an image-mode virtual disk. When you create an image-mode virtual disk, a direct mapping is made between extents that are on the managed disk and extents that are on the virtual disk. The managed disk is not virtualized. In other words, the logical block address (LBA) x on the managed disk is the same as LBA x on the virtual disk.

When you create an image-mode VDisk, you must assign it to a managed disk group. An image-mode VDisk must be at least one extent in size. In other words, the minimum size of an image-mode VDisk is the extent size of the MDisk group to which it is assigned.

The extents are managed in the same way as other VDIs. When the extents have been created, you can move the data onto other MDisks that are in the group without losing access to the data. After you move one or more extents, the virtual disk becomes a real virtualized disk, and the mode of the managed disk changes from image to managed.

Attention: If you add an MDisk to an MDisk group as a managed disk, any data on the MDisk will be lost. Ensure that you create image-mode VDIs from the MDisks that contain data before you start adding any MDisks to groups.

MDisks that contain existing data have an initial mode of unmanaged, and the cluster cannot determine whether it contains partitions or data.

A virtual disk can have one of three states. Table 13 describes the different states of a virtual disk:

Table 13. Virtual disk status

Status	Description
Online	The virtual disk is online and available if both nodes in the I/O group can access the virtual disk. A single node will only be able to access a VDisk if it can access all the MDisks in the MDisk group that are associated with the VDisk.
Offline	The VDisk is offline and unavailable if both nodes in the I/O group are missing or none of the nodes in the I/O group that are present can access the VDisk.
Degraded	The status of the virtual disk is degraded if one node in the I/O group is online and the other node is either missing or cannot access the virtual disk.

You can also use more sophisticated extent allocation policies to create VDIs. When you create a striped virtual disk, you can specify the same managed disk more than once in the list of MDisks that are used as the stripe set. This is useful if you have a managed disk group in which not all the MDisks are of the same capacity. For example, if you have a managed disk group that has two 18 GB MDisks and two 36 GB MDisks, you can create a striped virtual disk by specifying each of the 36 GB MDisks twice in the stripe set so that two-thirds of the storage is allocated from the 36 GB disks.

If you delete a virtual disk, you destroy access to the data that is on the virtual disk. The extents that were used in the virtual disk are returned to the pool of free extents that is in the managed disk group. The deletion might fail if the virtual disk is still mapped to hosts. The deletion might also fail if the virtual disk is still part of a FlashCopy or a Metro Mirror mapping. If the deletion fails, you can specify the force-delete flag to delete both the virtual disk and the associated mappings to hosts. Forcing the deletion will also delete the copy services relationship and mappings.

Related concepts

“Virtualization” on page 1

Virtualization is a concept that applies to many areas of the information technology industry.

Virtual disk-to-host mapping

Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDIs) within the SAN Volume Controller.

Virtual disk-to-host mapping is similar in concept to logical unit number (LUN) mapping or masking. LUN mapping is the process of controlling which hosts have

access to specific logical units (LUs) within the disk controllers. LUN mapping is typically done at the disk controller level. Virtual disk-to-host mapping is done at the SAN Volume Controller level.

Application servers can only access VDisks that have been made accessible to them. The SAN Volume Controller detects the fibre-channel ports that are connected to the SAN. These correspond to the host bus adapter (HBA) worldwide port names (WWPNs) that are present in the application servers. The SAN Volume Controller enables you to create logical hosts that group together WWPNs belonging to a single application server. VDisks can then be mapped to a host. The act of mapping a virtual disk to a host makes the virtual disk accessible to the WWPNs in that host, and therefore the application server itself.

VDisks and host mappings

The SAN concept known as LUN masking usually requires device driver software in each host. The device driver software masks the LUNs as instructed by the user. After the masking has been done, only some disks are visible to the operating system. The SAN Volume Controller performs a similar function, but, by default, it presents to the host only those VDisks that are mapped to that host. You must therefore map the VDisks to the hosts that are to access those VDisks.

Each host mapping associates a virtual disk with a host object and allows all HBA ports in the host object to access the virtual disk. You can map a virtual disk to multiple host objects. When a mapping is created, multiple paths might exist across the SAN fabric from the hosts to the SAN Volume Controllers that are presenting the virtual disk. Most operating systems present each path to a virtual disk as a separate storage device. The SAN Volume Controller, therefore, needs the IBM Subsystem Device Driver (SDD) software to be running on the host. This software handles the many paths that are available to the virtual disk and presents a single storage device to the operating system.

When you map a virtual disk to a host, you can optionally specify a SCSI ID for the virtual disk. This ID controls the sequence in which the VDisks are presented to the host. Take care when you specify a SCSI ID, because some device drivers stop looking for disks if they find an empty slot. For example, if you present three VDisks to the host, and those VDisks have SCSI IDs of 0, 1, and 3, the virtual disk that has an ID of 3 might not be found because no disk is mapped with an ID of 2. The cluster automatically assigns the next available SCSI ID if none is entered.

Figure 15 on page 63 and Figure 16 on page 63 show two VDisks, and the mappings that exist between the host objects and these VDisks.

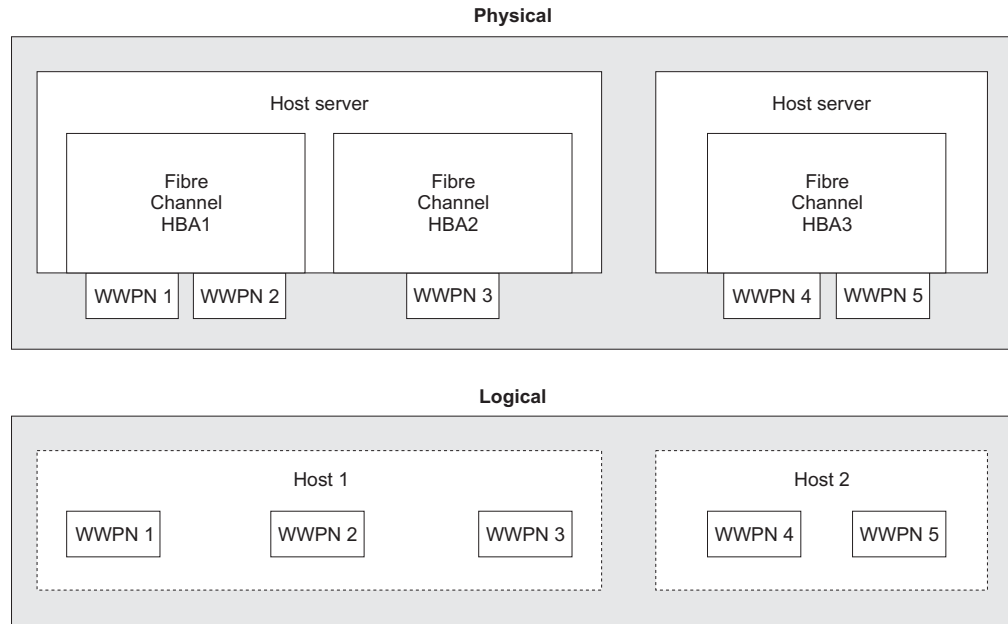


Figure 15. Hosts, WWPNs, and VDisks

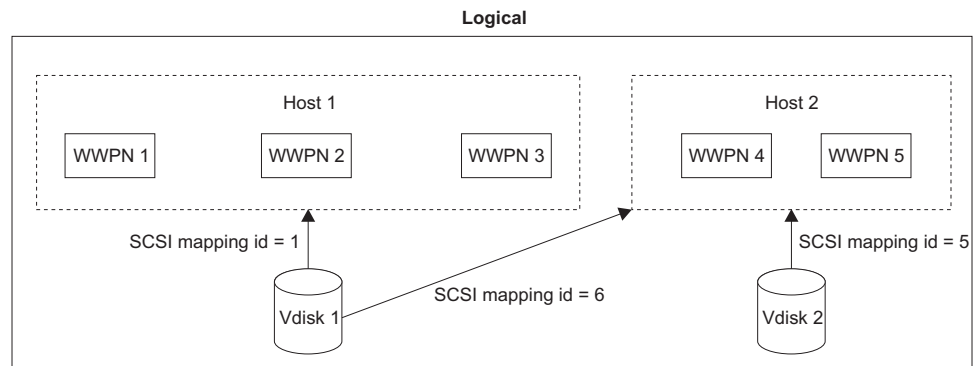


Figure 16. Hosts, WWPNs, VDisks and SCSI mappings

Host objects

A host system is an open-systems computer that is connected to the switch through a fibre-channel interface.

Creating a host in a cluster results in the creation of a logical host object. A logical host object has one or more worldwide port names (WWPNs) that are assigned to it. Generally, a logical host object is associated with a physical host system. However, a single logical host object can have WWPNs from multiple physical host systems that are assigned to it.

A host object is a logical object that groups one or more worldwide port names (WWPNs) of the host bus adapters (HBAs) that the cluster has detected on the SAN. A typical configuration has one host object for each host that is attached to the SAN. If, however, a cluster of hosts is going to access the same storage, you can add HBA ports from several hosts into the one host object to make a simpler configuration.

The cluster does not automatically present VDisks on the fibre. You must map each virtual disk to a particular set of ports to enable the virtual disk to be accessed through those ports. The mapping is made between a host object and a virtual disk.

When you create a new host object, the configuration interfaces provide a list of unconfigured WWPNs. These WWPNs represent the fibre channel ports that the cluster has detected.

The cluster can detect only ports that are logged into the fabric. Some HBA device drivers do not let the ports remain logged in if no disks are visible on the fabric. This condition causes a problem when you want to create a host because, at this time, no VDisks are mapped to the host. The configuration interface provides a method by which you can manually enter port names under this condition.

Attention: You must not include a node port in a host object.

A port can be added to only one host object. When a port has been added to a host object, that port becomes a configured WWPN, and is not included in the list of ports that are available to be added to other hosts.

Node login counts

The number of nodes that can see each port is reported on a per node basis and is known as the node login count. If the count is less than the number of nodes in the cluster, then there is a fabric problem, and not all nodes can see the port.

Chapter 6. Planning for configuring the SAN Volume Controller

Before you configure the SAN Volume Controller, you must complete these planning tasks.

Planning the clusters

Before you create a SAN Volume Controller cluster:

- Determine the number of clusters and the number of pairs of nodes. Each pair of nodes (the I/O group) is the container for one or more VDIs.
- Determine the number of hosts that will be used with the SAN Volume Controller. Hosts should be grouped by operating system and by type of host bus adapter (HBA).
- Determine the number of I/Os per second between the hosts and SAN Volume Controller nodes.

Planning the host groups

Host systems have access to specific logical units (LUs) within the disk controllers as a result of LUN masking. To plan a host group, gather the following information:

- List all of the worldwide port names (WWPNs) of the fibre-channel host bus adapter ports in the hosts.
- Determine the name to assign to the host or host group.
- Determine the VDIs to assign to the host.

Planning the managed disks

To plan the managed disks (MDisks), determine the logical or physical disks (logical units) in the backend storage.

Planning the managed disk groups

Before you create managed disk (MDisk) groups, determine the following factors:

- Determine the types of backend controllers in the system.
- If you want to create VDIs with the sequential policy, plan to create a separate MDisk group for these VDIs or ensure that you create these VDIs before creating VDIs with the striped policy.
- Plan to create MDisk groups for the backend controllers that provide the same level of performance or reliability, or both. For example, you can group all of the managed disks that are RAID 10 in one MDisk group and all of the MDisks that are RAID 5 in another group.

Planning the virtual disks

An individual virtual disk is a member of one managed disk group and one I/O group. The managed disk group defines which managed disks provide the backend storage that makes up the virtual disk. The I/O group defines which SAN Volume Controller nodes provide I/O access to the virtual disk. Determine the following information before creating a virtual disk:

- The name to assign to the virtual disk.
- The I/O group to which the virtual disk will be assigned.

- The managed disk group to which the virtual disk will be assigned.
- The capacity of the virtual disk.

Maximum configuration

Ensure that you are familiar with the maximum configurations of the SAN Volume Controller.

Table 14 shows the maximum configuration values to consider when you are planning the SAN Volume Controller installation.

Table 14. SAN Volume Controller maximum configuration values

Objects	Maximum number	Comments
Cluster Properties		
Nodes	8	Arranged as four I/O groups.
I/O groups	4	Each containing two nodes.
MDisk group	128	-- --
MDisks	4096	Represents an average of 64 per controller.
Object MDisks per MDisk group	128	-- --
MDisk size	2 TB	Defined by 32-bit LBA limit.
Addressability	2.1 PB	Maximum extent size 512 MB, arbitrary limit of 2^{22} extents in map.
LU size	2 TB	Defined by 32-bit LBA limit.
Concurrent SCSI tasks (commands) per node	10000	-- --
Concurrent commands per node	2500	Assumes a backend latency of 100 ms.
Concurrent commands per FC port	2048	-- --
SDD	512 SAN Volume Controller vpaths per host	<p>One vpath is created for each VDisk mapped to a host. Although the SAN Volume Controller permits 512 VDIs to be mapped to a host, the SDD limit can be exceeded by either:</p> <ul style="list-style-type: none"> • Creating two (or more) host objects for one physical host and mapping more than 512 VDIs to the host using the multiple host objects. • Creating two (or more) clusters and mapping more than 512 VDIs to the host using the multiple clusters. <p>Note: Both of these operations are unsupported.</p>
VDIs per MDisk Group		Cluster limit applies.
Front-end Properties		

Table 14. SAN Volume Controller maximum configuration values (continued)

Objects	Maximum number	Comments
SAN ports	512	Maximum size of fabric, including all SAN Volume Controller nodes.
Fabrics	2	Dual fabric configurations.
Host IDs per cluster	256	A host ID is associated with a map table that associates SCSI LUNs with VDisks. It is also associated with one or more host worldwide port names.
Host ports per cluster	512	Up to 512 distinct host worldwide port names are recognized.
Host LUN size	2 TB	Defined by 32-bit LBA limit.
Virtual disks (VDisks)	4096	Includes managed-mode VDisks and image-mode VDisks.
VDisks per I/O group	1024	-- --
VDisks per host ID	512	The limit may be different based on host operating system.
VDisks-to-host mappings	20 000	-- --
Maximum persistent reservation keys	132 000	-- --
Back-end Properties		
Managed Disks (MDisks)	4096	Represents an average of 64 per worldwide node name.
Back-end Storage WWNNs	64	Maximum number of device fabric worldwide node name.
Back-end Storage WWPNS	256	16 ports per controller
LUs per back-end WWNN	4096	Maximum of 512 LUs presented for each worldwide node name.
WWNNs per subsystem	4	-- --
WWPNS per WWNN	16	The maximum number of ports per worldwide node name.
Preferred ports per subsystem	4	
Copy Services Properties		
Metro Mirror relationships per cluster	4096	-- --

Table 14. SAN Volume Controller maximum configuration values (continued)

Objects	Maximum number	Comments
Metro Mirror consistency groups	32	-- --
Metro Mirror VDisk per I/O group	16 TB	-- --
FlashCopy mappings	2048	Supports up to 512 FlashCopy mappings per consistency group.
FlashCopy consistency groups	128	-- --
FlashCopy VDisk per I/O group	16 TB	-- --

Configuration rules and requirements

Ensure that you understand the rules and requirements when configuring the SAN Volume Controller.

The following terms and definitions will guide you in understanding the rules and requirements:

Table 15. Configuration Terms and Definitions

Term	Definition
ISL hop	A hop on an Inter-Switch Link (ISL). With reference to all pairs of N-ports or end-nodes that are in a fabric, the number of ISL hops is the number of links that are crossed on the shortest route between the node pair whose nodes are farthest apart from each other. The distance is measured only in terms of the ISL links that are in the fabric.
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 ISLs or where more than one ISL is 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. A symmetrical network means that all initiators are connected at the same level and all the controllers are connected at the same level. The SAN Volume Controller makes this calculation difficult, because it puts its back-end traffic onto the same network, and this back-end traffic varies by workload. Therefore, the oversubscription that a 100% read hit gives is different from the oversubscription that 100% write-miss gives. If you have an oversubscription of 1 or less, the network is nonblocking.
Virtual SAN (VSAN)	A VSAN is a virtual storage area network (SAN).
Redundant SAN	A SAN configuration in which if any one component fails, connectivity between the devices that are in the SAN is maintained, possibly with degraded performance. Create a redundant SAN by splitting the SAN into two independent counterpart SANs.

Table 15. Configuration Terms and Definitions (continued)

Term	Definition
Counterpart SAN	A non-redundant portion of a redundant SAN. A counterpart SAN provides all the connectivity of the redundant SAN, but without the redundancy. The SAN Volume Controller is typically connected to a redundant SAN that is made out of two counterpart SANs.
Local fabric	The fabric that consists of those SAN components (switches and cables) that connect the components (nodes, hosts, and switches) of the local cluster. Because the SAN Volume Controller supports Metro Mirror, significant distances might exist between the components of the local cluster and those of the remote cluster.
Remote fabric	The fabric that consists of those SAN components (switches and cables) that connect the components (nodes, hosts, and switches) of the remote cluster. Because the SAN Volume Controller supports Metro Mirror, significant distances might exist between the components of the local cluster and those of the remote cluster.
Local/remote fabric interconnect	The SAN components that connect the local fabrics to the remote fabrics. There may be significant distances between the components in the local cluster and those in the remote cluster. These components might be single-mode optical fibres that are driven by Gigabit Interface Converters (GBICs), or they might be other, more advanced components, such as channel extenders
SAN Volume Controller fibre-channel port fan in	The number of hosts that can see any one port. Some controllers recommend that the number of hosts using each port be limited to prevent excessive queuing at that port. If the port fails, or the path to that port fails, the host might failover to another port, and the fan in requirements might be exceeded in this degraded mode.
Invalid configuration	In an invalid configuration, an attempted operation will fail and will generate an error code to indicate what caused it to become invalid.
Unsupported configuration	A configuration that might operate successfully, but for which IBM does not guarantee to be able to solve problems that might occur. Usually this type of configuration does not create an error log entry.
Valid configuration	A configuration that is neither invalid nor unsupported.
Degraded	A valid configuration that has had a failure, but continues to be neither invalid nor unsupported. Typically, a repair action is required to restore the degraded configuration to a valid configuration.
Fibre Channel Extender	A device for long distance communication connecting other SAN fabric components. Generally these may involve protocol conversion to ATM, IP or some other long distance communication protocol.
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 ISL's for one of the diagonals. The SAN Volume Controller does not support this configuration.

Configuration rules

SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

Some configurations do not work and are known as *invalid*. You can avoid creating invalid configurations if you follow the rules that are given in this section.

A SAN configuration that contains SAN Volume Controllers is valid if it observes *all* of the following rules. These rules are discussed in the following section.

Storage subsystems

Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

All SAN Volume Controller nodes of a cluster must be able to see the same set of storage subsystem ports on each device. Any operation that is in this mode in which two nodes do not see the same set of ports on the same device is degraded, and the system logs errors that request a repair action. This rule can have important effects on a storage subsystem such as FAS*StT*, which has exclusion rules that determine to which host bus adapter (HBA) WWNNs a storage partition can be mapped.

A configuration in which a SAN Volume Controller bridges a separate host device and a RAID array is not supported. Typical compatibility matrixes are shown in a document titled *Supported Hardware List* on the following Web page:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

The SAN Volume Controller clusters must not share its storage subsystem logical units (LUs) with hosts. A storage subsystem can be shared with a host under certain conditions as described in this topic.

Certain storage controllers can be configured to safely share resources between the SAN Volume Controller and direct attached hosts. This type of configuration is described as a split controller. To determine support for a particular storage controller (and the restrictions on this support) by the SAN Volume Controller, refer to the section of this guide that describes configuring SAN devices and switches. In all cases, it is critical that the controller and SAN is configured so that the SAN Volume Controller cannot access logical units (LUs) that a host or another SAN Volume Controller can also access. This can be arranged by controller logical unit number (LUN) mapping and masking. If this is not guaranteed, then data corruption can occur.

Where SAN Volume Controller supports a controller being split between a SAN Volume Controller and a host as described above, the SAN Volume Controller also supports configurations in which a controller is split between two SAN Volume Controller clusters. In all cases, it is critical that the controller and SAN is configured so that the SAN Volume Controller cannot access LUs that a host or another SAN Volume Controller can also access. This can be arranged by controller LUN mapping and masking. If this is not guaranteed then data corruption can occur. This configuration is not recommended because of the risk of data corruption.

One storage subsystem device should not be configured to present the same LU to more than one SAN Volume Controller cluster. This configuration is not supported and is very likely to cause undetected data loss or corruption.

The SAN Volume Controller must be configured to manage only LUNs that are presented by supported disk controller systems. Operation with other devices is not supported.




Figure 17. Disk controller system shared between SAN Volume Controller and a host

| It is also possible to split a host so that it accesses some of its LUNs through the
 | SAN Volume Controller and some directly. In this case, the multi-pathing software
 | used by the controller must be compatible with SAN Volume Controllers
 | multi-pathing software. Figure 18 on page 73 is a supported configuration because
 | the same pathing driver is used for both direct and VDisks.
 |




Figure 18. ESS LUs accessed directly with a SAN Volume Controller

|
 | In the case where the RAID controller uses multi-pathing software which is
 | compatible with SAN Volume Controllers multi-pathing software, (see Figure 19 on
 | page 74) it is possible to configure a system such that some LUNs are mapped
 | directly to the host and others are accessed through SAN Volume Controller. One
 | example would be ESS which uses the same multi-pathing driver as SAN Volume
 | Controller.
 |




Figure 19. FAST direct connection with a SAN Volume Controller on one host

Related concepts

“Cluster operation and quorum disks” on page 38

The cluster must contain at least half of its nodes to function.

“Managed disks” on page 54

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

Host bus adapters

Follow these configuration rules for host bus adapters (HBAs).

SAN Volume Controller nodes always contain two HBAs. Each HBA must present two ports. If an HBA fails, the configuration is still valid, and the node operates in degraded mode. If an HBA is physically removed from a SAN Volume Controller node, the configuration is not supported.

HBAs that are in dissimilar hosts or dissimilar HBAs that are in the same host must be in separate zones. For example, if you have an HP/UX[®] host and a Windows[®] 2000 server host, those hosts must be in separate zones. Here, *dissimilar* means that the hosts are running different operating systems or that they are different hardware platforms. Different levels of the same operating system are considered to

be similar. This requirement ensures that different SANs can operate with each other. A configuration that breaks this requirement is not supported.

The SAN Volume Controller must be configured to export virtual disks (VDisks) only to host fibre-channel ports that are on the supported HBAs. See the following Web site for specific firmware levels and the latest supported hardware:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Operation with other HBAs is not supported.

The number of paths from the SAN Volume Controller nodes to a host must not exceed eight. The maximum number of HBA ports must not exceed four (for example, no more than two 2-port HBAs or four 1-port HBAs). Each SAN Volume Controller node in an I/O group presents four images of a VDisk onto the SAN, and each host SAN attachment has up to four HBA ports. Therefore, with more simplified zoning, the number of paths can equal up to 32: 4 SAN Volume Controller ports x 2 nodes per I/O group x 4 HBA ports. If you want to restrict the number of paths to a host, zone the switches so that each HBA port is zoned with one SAN Volume Controller port for each node in the cluster. If a host has multiple HBA ports, zone each port to a different set of SAN Volume Controller ports to maximize performance and redundancy.

Nodes

Follow these configuration rules for nodes.

The SAN Volume Controller nodes must always be deployed in pairs. If a node fails or is removed from the configuration, the remaining node operates in a degraded mode, but the configuration is still valid.

A SAN Volume Controller node always contains two host bus adapters (HBAs), each of which presents two ports. If an HBA fails, this remains a valid configuration, and the node operates in degraded mode. If an HBA is physically removed from a SAN Volume Controller node, then the configuration is unsupported.

Each SAN Volume Controller node presents a virtual disk (VDisk) to the SAN through four ports. Since each VDisk is accessible from the two SAN Volume Controller nodes in an I/O group, this means that a host HBA may see up to eight paths to each logical unit (LU) presented by the SAN Volume Controller. The hosts must run a multi-pathing device driver to resolve this back to a single device.

Support for optical connections is based on the fabric rules that the manufacturers impose for the following connection methods:

- Node to a switch
- Host to a switch
- Backend to a switch
- Switch to an Inter-Switch Link (ISL)

The optical connections supported between node and switch, host and switch, back-end and switch, and switch ISL should be determined by the fabric rules imposed by the vendors of the components used to connect the cluster.

For connections between an SAN Volume Controller node and its switch, the following optical connections are supported:

- Short wave optical fibre
- Long wave optical fibre up to 10 KM

High-power Gigabit Interface Converters (GBICs) and longwave fibre connections beyond 10 KM are not supported.

To ensure cluster failover operations, all nodes in a cluster must be connected to the same IP subnet.

The number of paths through the network from the SAN Volume Controller to a host must not exceed eight. Configurations in which this number is exceeded are unsupported. The SAN Volume Controller has four ports per node with two nodes in an I/O group. Therefore, without any zoning, the number of paths to a VDisks is eight * the number of host ports.

The nodes must be connected to the uninterruptible power supply using the supplied cable which joins together the signal and power cables.

All nodes in the cluster should be connected to the same IP subnet to ensure cluster failover operation.

Power requirements

Note the power requirements for the SAN Volume Controller.

The uninterruptible power supply must be in the same rack that contains the SAN Volume Controller nodes that it supplies. The combination power and signal cable for connection between SAN Volume Controller and uninterruptible power supply units is two meters long. The SAN Volume Controller and uninterruptible power supply must connect with both the power and the signal cable to function correctly.

Fibre-channel switches

Follow these guidelines for configuring the fibre-channel switches that are supported on the SAN.

The SAN must contain only supported switches. The SAN Volume Controller supports specific IBM 2109, McData, and InRange switch models and the Cisco MDS 9000 switch and switches supported by the Cisco MDS 9000.

See the following Web site for specific firmware levels and the latest supported hardware:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Operation with other switches is not supported.

Ideally, the SAN would consist of two independent switches (or networks of switches) so that the SAN includes a redundant fabric, and has no single point of failure. If one SAN fabric fails, the configuration is in a degraded mode, but it is still valid. If the SAN contains only one fabric, it is still a valid configuration, but a failure of the fabric might cause a loss of access to data. Such a SAN, therefore, is considered a possible single point of failure.

Within an individual SAN fabric, switches must be from the same manufacturer with the specific exceptions for BladeCenter.

Note: Where a pair of counterpart fabrics (for example, Fabric A and Fabric B) are used to provide a redundant SAN, it is supported for Fabric A to be comprised of one vendors switches and Fabric B to use a different vendors switch so long as each individual fabric complies with the following:

- Different vendor switches may be mixed in a SAN Volume Controller configuration, provided that each fabric contains only switches from a single vendor. Thus, the two counterpart SANs can have different vendor switches.

SANs that are created from a mesh of switches are not supported.

Configurations with more than two SANs are not supported.

The fibre-channel SAN connections between SAN Volume Controller and the switches are optical fibre running at 2 Gb/s, however, SAN Volume Controller is also supported in 1 Gb/s fibre-channel fabrics. The fibre-channel switch will need to be zoned to permit the hosts to see the SAN Volume Controller nodes and the SAN Volume Controller nodes to see the RAID controllers. The SAN Volume Controller nodes within a cluster must be able to see each other.

On the fibre-channel SAN, the SAN Volume Controller nodes must always and only be connected to SAN switches. Each node must be connected to each of the counterpart SANs that are in the redundant fabric. Any operation that uses direct connections between host and node, or controller and node, is not supported.

On the fibre-channel SAN, back-end storage must always and only be connected to SAN switches. Multiple connections are permitted from the redundant controllers of the back-end storage, to improve data bandwidth performance. It is not necessary to have a connection between each redundant disk controller system of the back-end storage and each counterpart SAN. For example, in a FAST configuration in which the FAST contains two redundant controllers, only two controller mini-hubs are usually used. Controller A of the FAST is, therefore, connected to counterpart SAN A, and controller B of the of the FAST is connected to counterpart SAN B. Any operation that uses a direct connection between the host and the controller is not supported.

The connections between the switches and the SAN Volume Controllers can operate at 1 Gbps or at 2 Gbps and are made with optical fibre. All the ports for the SAN Volume Controllers that are in a single cluster, however, must run at one speed. Any operation that runs different speeds on the node-to-switch connections that are in a single cluster is not valid. Mixed speeds are permitted within the fabric. You can use lower speeds to extend distance or to make use of 1 Gb/s legacy components.

Attention: The default transfer rate in the SAN Volume Controller is 2 Gbps. If your environment is set up to use 1 Gbps switches, the switch rate must be set at the transfer rate.

Mixed speeds are permitted in the fabric. Lower speeds can be used to extend distances or to make use of 1 Gbps legacy components.

Each device has a fibre-channel connection and an Ethernet connection for configuration and error reporting. These connections are aggregated together through an Ethernet hub.

When attaching SAN Volume Controller to a SAN fabric containing core directors and edge switches, it is preferable to connect the SAN Volume Controller ports to the core directors and to connect the host ports to the edge switches. In such a fabric, the next priority for connection to the core directors is the storage controllers, leaving the host ports connected to the edge switches.

The switch configuration of a SAN Volume Controller SAN must observe the switch manufacturer's configuration rules. These rules might put restrictions on the switch configuration; for example, the switch manufacturer might not permit other manufacturer's switches to be in the SAN. Any operations that run outside the manufacturer's rules is not supported.

The switch must be configured so that the SAN Volume Controller nodes can see the back-end storage and the front-end HBAs. However, the front-end HBAs and the back-end storage must not be in the same zone. Any operation that runs outside these zoning rules is not supported.

All SAN Volume Controller nodes in a SAN Volume Controller cluster must be able to see the same set of back-end storage ports on each back-end controller. Operation in a mode where two nodes see a different set of ports on the same controller is degraded and the system will log errors requesting a repair action. This could occur if inappropriate zoning was applied to the fabric. It could also occur if inappropriate LUN masking is used. This rule has important implications for back-end storage, such as FAStT which imposes exclusive rules on which HBA WWN's a storage partition can be mapped to.

Because each SAN Volume Controller has four ports, the switches can be zoned so that a particular SAN Volume Controller port is used only for internode communication, for communication to the host, or for communication to back-end storage. Whatever the configuration, each SAN Volume Controller node must remain connected to the full SAN fabric. Zoning must not be used to split the SAN into two parts.

Fibre-channel switches and Inter-Switch Links

The local or remote fabric must not contain more than three Inter-Switch Links (ISLs) hops in each fabric. Any operation that uses more than three hops is not supported. When a local fabric is connected to a remote fabric for Metro Mirror, the hop count between a local node and a remote node must not exceed seven. Therefore, some hops can be used in a cascaded switch link between local and remote clusters if the internal ISL count of the local or remote cluster is less than three.

Note: You must use trunking when you use multiple ISLs in parallel.

The local and remote fabric interconnections must be only one ISL hop between a switch that is in the local fabric and a switch that is in a remote fabric. That is, it must be a single-mode fibre up to 10 KM (32 810 ft.) long. Any operation that uses other local or remote fabric interconnections is not supported.

When ISLs are used, each ISL oversubscription must not exceed six. Any operation that uses higher values is not supported.

With Inter-Switch Links between nodes in the same cluster, the ISLs are considered a single point of failure. This is illustrated in Figure 20 on page 79.




Figure 20. Fabric with Inter-Switch Links between nodes in a cluster

If Link 1 or Link 2 fails, the cluster communication does not fail.

If Link 3 or Link 4 fails, the cluster communication does not fail.

If ISL 1 or ISL 2 fails, the communication between Node A and Node B will fail for a period of time, and the node will not be recognized, even though there is still a connection between the nodes.

To ensure that a fibre-channel link failure does not cause nodes to fail when there are ISLs between nodes, it is necessary to use a redundant configuration. This is illustrated in Figure 21.

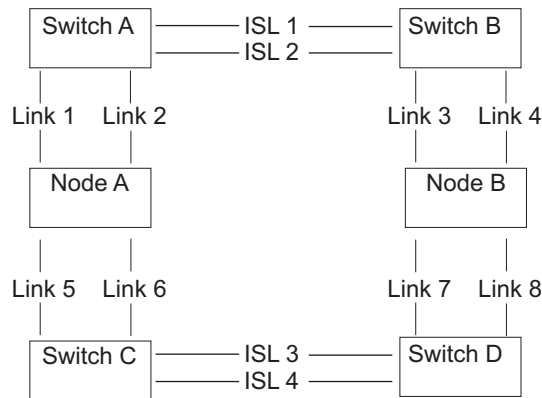


Figure 21. Fabric with Inter-Switch Links in a redundant configuration

With a redundant configuration, if any one of the links fails, then communication on the cluster will not fail.

SAN Volume Controller in a SAN with director class switches

When a large number of RAID controllers and test hosts are to be connected to a SAN Volume Controller cluster, it is possible that director class switches will be employed within the SAN. Director class switches can offer internal redundancy so it may be possible to replace a SAN that employs multiple switches to provide two redundant SAN fabrics with a single director class switch. While this may give adequate SAN network redundancy, this does not protect against physical damage (for example, flood or fire), which would destroy the entire function in the director class switch. It may be more attractive therefore to use tiered networks of smaller switches or a core-edge topology, with multiple switches in the core. The redundancy in the network can then be physically distributed across a wide area protecting against physical damage.

Configuration requirements

You *must* perform these steps before you configure the SAN Volume Controller.

1. Your IBM service representative must have installed the SAN Volume Controller.
2. Install and configure your disk controller systems and create the RAID resources that you intend to virtualize. To prevent loss of data, virtualize only those RAID that provide some kind of redundancy, that is, RAID 1, RAID 10, RAID 0+1, or RAID 5. Do *not* use RAID 0 because a single physical disk failure might cause the failure of many virtual disks. RAID 0, like other types of RAID offers cost-effective performance by using available capacity through data striping. However, RAID 0 does not provide a parity disk drive for redundancy (RAID 5) or mirroring (RAID 10).

When creating RAID with parity protection (for example, RAID 5), consider how many component disks to use in each array. The more disks you use, the fewer disks you need to provide availability for the same total capacity (one per array). However, if you use more disks, it will take longer to rebuild a replacement disk after a disk failure. If a second disk failure occurs during the rebuild period, all data on the array is lost. More data is affected by a disk failure for a larger number of member disks resulting in reduced performance while rebuilding onto a hot spare and more data being exposed if a second disk fails before the rebuild has completed. The smaller the number of disks, the more likely it is that write operations span an entire stripe (strip size x number of members minus 1). In this case, write performance is improved because the disk write operations do not have to be preceded by disk reads. The number of disk drives required to provide availability might be unacceptable if the arrays are too small.

When in doubt, create arrays with between six and eight member disks.

If reasonably small RAID arrays are used, it is easier to extend an MDisk group by adding a new RAID array of the same type. Construct multiple RAID devices of the same type, when possible.

When creating RAID with mirroring, the number of component disks in each array does not affect redundancy or performance.

Most back-end disk controller systems enable RAID to be divided up into more than one SCSI logical unit (LU). When configuring new storage for use with the SAN Volume Controller, you do not need to divide up the array. New storage should be presented as one SCSI LU. This will give a one-to-one relationship between MDisks and RAID.

Attention: Losing an array in an MDisk group can result in the loss of access to *all* MDisks in that group.

3. Install and configure your switches to create the zones that the SAN Volume Controller requires. One zone must contain all the disk controller systems and the SAN Volume Controller nodes. For hosts with more than one port, use switch zoning to ensure that each host fibre-channel port is zoned to exactly one fibre-channel port of each SAN Volume Controller node in the cluster. Set up a zone on each fibre-channel switch that includes the master console and all of the SAN Volume Controller ports that are connected to that switch.
4. If you want the SAN Volume Controller to export redundant paths to disks, you must install the Subsystem Device Driver (SDD) on all of the hosts that are connected to the SAN Volume Controller. Otherwise, you will not be able to use the redundancy inherent in the configuration. Install the SDD from the following Web site:

<http://www-1.ibm.com/server/storage/support/software/sdd.html>

Version 1.4.x.x or later is required.

5. Install and configure the SAN Volume Controller master console (see the *IBM TotalStorage Master Console Installation and User's Guide*). The communication between the master console and the SAN Volume Controller runs under a client-server network application called Secure Shell (SSH). Each SAN Volume Controller cluster is equipped with SSH Server software and the master console comes to you equipped with the SSH Client software called PuTTY. You will need to configure the SSH client key pair using PuTTY on the master console. Once you have installed your master console, you can configure and administer the SAN Volume Controller using a graphical interface or a command-line interface.

a. You can configure the SAN Volume Controller using the SAN Volume Controller Console Web-based application that is preinstalled on the master console.

Note: You can also install the master console on another machine (which you provide) using the CD-ROM provided with the master console.

b. You can configure the SAN Volume Controller using the command-line interface (CLI) commands.

c. You can install an SSH client if you only want to use the CLI commands. If you want to use the CLI from a host other than the master console, ensure that the host has an SSH client installed on it.

Note:

- 1) AIX comes with an installed SSH client.
- 2) Linux comes with an installed SSH client.
- 3) PuTTY is recommended for Windows.

When you and the IBM service representative have completed the initial preparation steps, you must perform the following steps:

1. Add nodes to the cluster and set up the cluster properties.
2. Create managed disk groups from the managed disks to make pools of storage from which you can create virtual disks.
3. Create host objects from the HBA fibre-channel ports to which you can map virtual disks.
4. Create virtual disks from the capacity that is available in your managed disk groups.
5. Map the virtual disks to the host objects to make the disks available to the hosts, as required.
6. Optionally, create Copy Services (FlashCopy and Metro Mirror) objects, as required.

Related concepts

“Managed disk groups” on page 56

An *MDisk group* is a collection of MDiskS that jointly contain all the data for a specified set of virtual disks (VDiskS).

Related reference

“Fibre-channel switches” on page 76

Follow these guidelines for configuring the fibre-channel switches that are supported on the SAN.

Chapter 7. SAN Volume Controller supported environment

The IBM Web site provides up-to-date information about the supported environment for the SAN Volume Controller.

This includes:

- Host attachments
- Physical disk storage systems
- Host bus adapters
- Switches

See the following Web site for specific firmware levels and the latest supported hardware:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Supported host attachments

The IBM Web site provides up-to-date information about the supported host attachment operating systems.

For a list of supported host attachment operating systems, see the SAN Volume Controller Web site at:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

The SAN Volume Controller provides heterogeneous host attachments so that you can consolidate storage capacity and workloads for open systems hosts. The SAN Volume Controller supports a maximum of 256 separate hosts and a maximum of 512 host fibre-channel ports, identified by their worldwide port numbers (WWPNs).

Hosts are attached to the SAN Volume Controller using a switched, fibre-channel fabric.

Supported storage subsystems

The IBM Web site provides up-to-date information about the supported physical disk storage systems.

For a list of supported storage systems, see the SAN Volume Controller Web site:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>.

Supported fibre-channel host bus adapters

The IBM Web site provides up-to-date information about the supported host bus adapters.

Ensure that host bus adapters (HBAs) are at or above the minimum requirements.

For a list of supported HBAs, see the following Web site for specific firmware levels and the latest supported hardware:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Supported switches

The IBM Web site provides up-to-date information about the supported fibre-channel switches.

Ensure that switches are at or above the minimum requirements.

The SAN must contain only supported switches.

See the following Web site for the latest models and firmware levels:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Operation with other switches is not supported.

Supported fibre-channel extenders

The SAN Volume Controller supports the CNT UltraNet Edge Storage Router, the Cisco MDS 9506 and 9509 Director, and Cisco 9516 Fabric Switches to support synchronous copy services.

The maximum one-way latency supported is 10 milliseconds when using Brocade fabric, and 34 milliseconds when using McData fabric. The relationship between latency and distance is dependent on the network and number of hops. The distance is approximately 100-150 kilometers per millisecond.

Note: Performance of copy services degrades as the distance increases.

The supported hardware for the SAN Volume Controller frequently changes. See the following Web site for the latest supported hardware:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the SAN Volume Controller master console:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen readers have been tested: JAWS v4.5 and IBM Home Page Reader v3.0.
- You can operate all features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the SAN Volume Controller Console and help system from the keyboard by using the following key combinations:

- To traverse to the next link, button, or topic, press Tab inside a frame (page).
- To expand or collapse a tree node, press → or ←, respectively.
- To move to the next topic node, press V or Tab.
- To move to the previous topic node, press ^ or Shift+Tab.
- To scroll all the way up or down, press Home or End, respectively.
- To go back, press Alt+←.
- To go forward, press Alt+→.
- To go to the next frame, press Ctrl+Tab.
- To move to the previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.
- To select, press Enter.

Accessing the publications

You can view the publications for the SAN Volume Controller in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided on a CD that is packaged with the product or you can access them at the following Web site:

<http://www-1.ibm.com/servers/storage/support/virtual/2145.html>

Related reference

“SAN Volume Controller library and related publications” on page xi
A list of other publications that are related to this product are provided to you for your reference.

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 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.

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:

- AIX
- e (logo)
- Enterprise Storage Server
- FlashCopy
- IBM
- Tivoli
- TotalStorage
- xSeries

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.

Definitions of notices

Ensure that you understand the typographic conventions that are used in this publication to indicate special notices.

The following notices are used throughout this library to convey the following specific meanings:

Note: These notices provide important tips, guidance, or advice.

Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage could occur.

CAUTION:

These notices indicate situations that can be potentially hazardous to you. A caution notice precedes the description of a potentially hazardous procedural step or situation.

DANGER

These notices indicate situations that can be potentially lethal or extremely hazardous to you. A danger notice precedes the description of a potentially lethal or extremely hazardous procedural step or situation.

Glossary

Ensure you are familiar with the list of terms and their definitions used in this guide.

A

asymmetric virtualization

A virtualization technique in which the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and locking tables while the storage devices contain only data. See also *symmetric virtualization*

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*.

B

blade One component in a system that is designed to accept some number of components (blades). Blades could be individual servers that plug into a multiprocessing system or individual port cards that add connectivity to a switch. A blade is typically a hot-swappable hardware device.

block A unit of data storage on a disk drive.

block virtualization

The act of applying virtualization to one or more block-based (storage) services for the purpose of providing a new aggregated, higher-level, richer, simpler, or secure block service to clients. Block virtualization functions can be nested. A disk drive, RAID system, or volume manager all perform some form of block-address to (different) block-address mapping or aggregation. See also *virtualization*.

C

cache A high-speed memory or storage device used to reduce the effective time required to read data from or write data to lower-speed memory or a device. Read cache holds data in anticipation that it will be requested by a client. Write cache holds data written by a client until it can be safely stored on more permanent storage media such as disk or tape.

cascading

The process of connecting two or more fibre-channel hubs or switches together to increase the number of ports or extend distances.

cluster

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

CIM See *Common Information Model*.

CLI See *command line interface*.

command line-interface (CLI)

A type of computer interface in which the input command is a string of text characters.

Common Information Model (CIM)

A set of standards developed by the Distributed Management Task Force (DMTF). CIM provides a conceptual framework for storage management

and an open approach to the design and implementation of storage systems, applications, databases, networks, and devices.

connected

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

consistency group

A group of copy relationships between virtual disks that are managed as a single entity.

consistent copy

In a Global Mirror relationship, a copy of a secondary virtual disk (VDisk) that is identical to the primary VDisk from the viewpoint of a host system, even if a power failure occurred while I/O activity was in progress.

copied

In a FlashCopy relationship, a state that indicates that a copy has been started after the copy relationship was created. The copy process is complete and the target disk has no further dependence on the source disk.

Copy Services

In the SAN Volume Controller, the two services that enable you to copy virtual disks (VDisks): FlashCopy and Global Mirror.

copying

A status condition that describes the state of a pair of virtual disks (VDisks) that have a copy relationship. The copy process has been started but the two virtual disks are not yet synchronized.

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.

D

data migration

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

destage

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

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).

E

error code

A value that identifies an error condition.

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.

F

fabric In fibre-channel technology, a routing structure, such as a switch, that receives addressed information and routes it to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are described as cascading. See also *cascading*.

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.

fibre channel

A technology for transmitting data between computer devices at a data rate of up to 4 Gbps. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives.

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*.

H

HBA See *host bus adapter*.

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).

hub A communications infrastructure device to which nodes on a multi-point bus or loop are physically connected. Commonly used in Ethernet and fibre-channel networks to improve the manageability of physical cables. Hubs maintain the logical loop topology of the network of which they are a part, while creating a “hub and spoke” physical star layout. Unlike switches,

hubs do not aggregate bandwidth. Hubs typically support the addition or removal of nodes from the bus while it is operating. (S) Contrast with *switch*.

I

IBM TotalStorage Multipath Subsystem Device Driver (SDD)

An IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

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.

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*.

I/O group

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

inconsistent

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

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.

Internet Protocol (IP)

In the Internet suite of protocols, a connectionless protocol that routes data through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

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.

Inter-Switch Link (ISL)

A protocol for interconnecting multiple routers and switches in a storage area network.

ISL See *Inter-Switch Link*.

ISL hop

Considering all pairs of node ports (N-ports) in a fabric and measuring distance only in terms of Inter-Switch Links (ISLs) in the fabric, the number of ISLs traversed is the number of ISL hops on the shortest route between the pair of nodes that are farthest apart in the fabric.

IP See *Internet Protocol*.

I/O See *input/output*.

J

JBOD (just a bunch of disks)

IBM definition: See *non-RAID*. HP definition: A group of single-device logical units not configured into any other container type.

L

line card

See *blade*.

local fabric

In SAN Volume Controller, those storage area network (SAN) components (such as switches and cables) that connect the components (nodes, hosts, switches) of the local cluster together.

logical unit (LU)

An entity to which Small Computer System Interface (SCSI) commands are addressed, such as a virtual disk (VDisk) or managed disk (MDisk).

logical unit number (LUN)

The SCSI identifier of a logical unit within a target. (S)

LU See *logical unit*.

LUN See *logical unit number*.

M

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*.

Management Information Base (MIB)

Simple Network Management Protocol (SNMP) units of managed information that specifically describe an aspect of a system, such as the system name, hardware number, or communications configuration. A collection of related MIB objects is defined as a MIB.

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*.

migration

See *data migration*.

N

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

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.

non-RAID

Disks that are not in a redundant array of independent disks (RAID). IBM definition: Disks that are not in a redundant array of independent disks (RAID). HP definition: See *JBOD*.

O

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.

P**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 *T₀ copy*.

port ID

An identifier associated with a port.

primary virtual disk

In a Global Mirror relationship, the target of write operations issued by the host application.

PuTTY

A free implementation of Telnet and SSH for Windows 32-bit platforms

Q**quorum disk**

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

R

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

RAID See *redundant array of independent disks*.

redundant array of independent disks

A collection of two or more disk drives that present the image of a single disk drive to the system. In the event of a single device failure, the data can be read or regenerated from the other disk drives in the array.

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*.

relationship

In 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*.

Global Mirror

In SAN Volume Controller, a copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk designated in the relationship.

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.

S

SAN See *storage area network*.

SDD See *IBM TotalStorage Multipath Subsystem Device Driver (SDD)*.

secondary virtual disk

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

Simple Network Management Protocol (SNMP)

In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application-layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

SNMP See *Simple Network Management Protocol*.

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)

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*.

symmetric virtualization

A virtualization technique in which the physical storage in the form of Redundant Array of Independent Disks (RAID) is split into smaller chunks of storage known as *extents*. These extents are then concatenated, using various policies, to make virtual disks (VDisks). See also *asymmetric virtualization*.

synchronized

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

U

uninterruptible power supply

A device connected between a computer and its power source that protects the computer against blackouts, brownouts, and power surges. The uninterruptible power supply contains a power sensor to monitor the supply and a battery to provide power until an orderly shutdown of the system can be performed.

unconfigured mode

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

V

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.

virtual storage area network (VSAN)

A fabric within the SAN.

virtualization

In the storage industry, a concept in which a pool of storage is created that contains several disk subsystems. The subsystems can be from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them.

virtualized storage

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

VLUN See *virtual disk*.

VSAN See *virtual storage area network*.

W

worldwide node name (WWNN)

An identifier for an object that is globally unique. WWNNs are used by Fibre Channel and other standards.

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

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.

Index

A

- accessibility
 - keyboard 85
 - shortcut keys 85
- adapters
 - fibre channel 83

C

- cable connection table
 - example 25
- charts and tables 21
 - cable connection table 25
 - configuration data table 27
 - hardware location chart 21, 22, 23
- Cisco Systems
 - MDS 9000 Caching Services Module 5
 - MDS 9000 switch 5
- cluster state 38
- clusters
 - operation 38
 - operation over long distances 35
 - overview 37
- configuration
 - maximum sizes 66
 - rules 68
- configuration requirements 80
- configuring
 - switches 76
- connections 20
- consistency group, Metro Mirror 49
- consistency groups, FlashCopy 46
- conventions
 - emphasis in text xi
- copy services
 - overview 43
- country power cables 16, 18

D

- disk controllers
 - overview 53

E

- emphasis in text xi

F

- FlashCopy
 - consistency groups 46
 - mappings 44
 - overview 44

H

- HBAs (host bus adapters)
 - configuration 74
- host bus adapters (HBAs)
 - configuration 74
- hosts 83
 - overview 63

I

- I/O groups 39
- IBM TotalStorage Metro Mirror for SAN Volume Controller 49
- information
 - center xi
- installation
 - planning 15, 21, 29

K

- keyboard 85
- shortcut keys 85

M

- Metro Mirror
 - overview 48, 49
 - zoning considerations 33
- migration 43

N

- node status 36
- nodes
 - configuration 75
- notices 89
 - legal 87

O

- object descriptions 51
- operating over long distances 34
- ordering publications xiii
- overview
 - disk controllers 42
 - zoning 30

P

- physical characteristics
 - uninterruptible power supply 19
- planning
 - configuration 65
 - installation 15, 21, 29
- ports 20

- power
 - SAN Volume Controller requirements 15
- power cables 5115
 - country 16
 - region 16
- power cables 5125
 - country 18
 - region 18
- power requirements 76
- publications
 - ordering xiii

R

- related information xi
- requirements
 - ac voltage 15
 - electrical 15
 - power 15

S

- safety
 - caution notices 89
 - danger notices 89
- SAN Volume Controller
 - air temperature 15
 - dimensions and weight 15
 - heat output 15
 - humidity 15
 - overview 5
 - product characteristics 15
 - specifications 15
 - weight and dimensions 15
- SANs (storage area networks) 29
- shortcut keys 85
- site requirements
 - connections 20
 - ports 20
- status
 - of cluster 38
 - of node 36
- storage
 - devices
 - supported 83
- storage area networks (SAN) 29
- support
 - Web sites xii
- switches
 - operating over long distances 34
 - supported 84

T

- text emphasis xi
- trademarks 88

U

- uninterruptible power supply
 - environment 19
 - overview 9, 40

V

- virtual disk-to-host mapping
 - description 61
- virtual disks (VDisks)
 - modes 43
- virtualization
 - asymmetric 3
 - overview 1
 - symmetric 4

W

- Web sites xii

Z

- zoning
 - considerations for Metro Mirror 33
 - overview 30

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

IBM TotalStorage SAN Volume Controller
Planning Guide
Version 2.1.0

Publication No. GA22-1052-04

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

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.

Name

Address

Company or Organization

Phone No.



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
USA 85775-4401



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

GA22-1052-04



Spine information:



IBM TotalStorage SAN Volume
Controller

SAN Volume Controller Planning Guide

Version 2.1.0