

IBM System Storage SAN Volume Controller



Planning Guide

Version 4.3.0

IBM System Storage SAN Volume Controller



Planning Guide

Version 4.3.0

Note:

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

This edition applies to the IBM System Storage SAN Volume Controller, release 4.3.0, and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces GA32-0551-02.

© **Copyright International Business Machines Corporation 2003, 2008. 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 GA32-0551-03 SAN

 Volume Controller Planning Guide ix

 Summary of changes for GA32-0551-02 SAN

 Volume Controller Planning Guide x

Emphasis x

SAN Volume Controller library and related

publications xi

Related Web sites xv

How to order IBM publications. xv

How to send your comments xv

**Chapter 1. SAN Volume Controller
overview 1**

SAN Volume Controller objects 2

SAN Volume Controller cluster high availability 3

SAN Volume Controller operating environment. 4

2145-1U uninterruptible power supply. 6

Redundant ac power switch 6

IBM System Storage Productivity Center 7

 Assist On-site and remote service 8

 Secure Shell protocol through PuTTY 9

Data and event notifications 9

Clusters 10

 Cluster state 10

 Cluster operation and quorum disks 11

I/O groups and uninterruptible power supply units 11

Disk controllers 13

Virtual Disk Mirroring feature 14

Space-Efficient Virtual Disk feature 14

Data migration 15

 Image mode virtual disk migration 16

Cluster configuration backup functions 16

Copy Services features. 16

 FlashCopy. 16

 Metro Mirror and Global Mirror 17

**Chapter 2. SAN Volume Controller
physical installation planning 19**

SAN Volume Controller 2145-8G4 environment

requirements 19

Requirements for the uninterruptible power supply

environment 21

 2145-1U uninterruptible power supply

 environment 21

 Power cables for the 2145-1U uninterruptible
 power supply 22

Preparing your redundant ac power environment 24

Connections 25

Preparing your IBM System Storage Productivity

Center environment 25

Web browser requirements to access the SAN

Volume Controller Console 25

Physical configuration planning of a system with

SAN Volume Controller 26

 Requirements and guidelines for completing the
 hardware location chart 26

 Requirements for completing the cable
 connection table 27

 Configuration data table guidelines 28

 Requirements for completing the redundant ac
 power switch connection chart 29

 Cabling of redundant ac power switch (example) 30

**Chapter 3. SAN fabric and zoning
overview 33**

SAN fabric overview 33

SAN fabric configuration 34

SAN hardware configuration 35

Zoning guidelines 38

Zoning examples 41

Zoning considerations for Metro Mirror and Global

Mirror 42

Switch operations over long distances 43

**Chapter 4. Configuration planning for
the SAN Volume Controller 45**

SAN Volume Controller maximum configuration 46

Configuration rules. 46

 Storage subsystem configuration rules 46

 Host bus adapter configuration rules 50

 Node configuration rules 51

Accessibility 55

Notices 57

Trademarks 59

Definitions of notices 59

Glossary 61

Index 85

Figures

	1.	Overview of the IBM System Storage Productivity Center	8		
	2.	Cluster, nodes, and cluster state.	11		
	3.	I/O group and uninterruptible power supply	12		
	4.	A four-node SAN Volume Controller cluster with the redundant ac power switch feature	31		
	5.	Example of a SAN Volume Controller cluster in a fabric	33		
	6.	Fabric with ISL between nodes in a cluster	37		
	7.	Fabric with ISL in a redundant configuration	38		
	8.	Disk controller system shared between SAN Volume Controller node and a host	48		
	9.	IBM System Storage DS8000 LUs accessed directly with a SAN Volume Controller node	49		
	10.	IBM DS4000 direct connection with a SAN Volume Controller node on one host	50		

Tables

- | | | | |
|--|----|--|----|
| 1. Configuration terms and definitions | 34 | 3. Six hosts and their ports | 42 |
| 2. Four hosts and their ports. | 41 | | |

About this guide

This publication introduces the IBM System Storage 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 System Storage 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 GA32-0551-03 SAN Volume Controller Planning Guide

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

New information

This topic describes the changes to this guide since the previous edition, GA32-0551-02. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Introduction of the Virtual Disk Mirroring feature.
- Introduction of the Space-Efficient Virtual Disk feature.
- Introduction of IBM System Storage Productivity Center.
- Implementation of Internet Protocol Version 6 (IPv6) on SAN Volume Controller in addition to the IPv4 standard currently in use.

Changed information

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

- Information about IBM System Storage Productivity Center was added.
- Information about connections was updated to include IPv6 support.
- Several topics within the manual were reorganized.

Removed Information

This section lists information that was removed from this book.

- Information about master console was removed from this guide. The information remains in an appendix in the *IBM System Storage SAN Volume Controller: Hardware Installation Guide*. It also is included in the *IBM System Storage SAN Volume Controller: Software Installation and Configuration Guide*.

Summary of changes for GA32-0551-02 SAN Volume Controller Planning Guide

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

New information

This topic describes the changes to this guide since the previous edition, GA32-0551-01. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- An example for cabling redundant ac power switch units.
- Confirmation that the master console and the SAN Volume Controller now support the Microsoft Windows Internet Explorer version 7.0 Web browser.

Changed information

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

- Added information about sending and receiving Call Home and Inventory information e-mails.
- Introduced the Incremental FlashCopy feature.
- Revised the cable numbers for the redundant ac power switch and the 2145-1U uninterruptible power supply and the outlet type for the power distribution unit.
- Described the longer distance now supported between the cluster and host or the cluster and the storage controller when using short or long wave optical fiber connections and SAN fabric switches.

Removed Information

This section lists information that was removed from this book.

- Moved information about SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 to Appendix A in the *IBM System Storage SAN Volume Controller: Hardware Installation Guide*.
- Moved information about SAN Volume Controller 2145-4F2 to Appendix B in the *IBM System Storage SAN Volume Controller: Hardware Installation Guide*.
- Moved information about the 2145 uninterruptible power supply to Appendix B in the *IBM System Storage SAN Volume Controller: Hardware Installation Guide*.

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 System Storage 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) from the following Web site:

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

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller: CIM Agent Developer's Reference</i>	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	SC26-7904
<i>IBM System Storage 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).	SC26-7903
<i>IBM System Storage SAN Volume Controller: Software Installation and Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller.	SC23-6628
<i>IBM System Storage SAN Volume Controller: Host Attachment Guide</i>	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7905
<i>IBM System Storage SAN Volume Controller: Hardware Installation Guide</i>	This guide includes the instructions that the IBM service representative uses to install the SAN Volume Controller hardware.	GC27-2132

Title	Description	Order number
<i>IBM System Storage 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.	GA32-0551
<i>IBM System Storage SAN Volume Controller: Service Guide</i>	This guide includes the instructions that the IBM service representative uses to service the SAN Volume Controller.	GC26-7901
<i>IBM Systems Safety Notices</i>	This guide contains translated caution and danger statements. Each caution and danger statement in the SAN Volume Controller documentation has a number that you can use to locate the corresponding statement in your language in the <i>IBM Systems Safety Notices</i> document.	G229-9054

Other IBM publications

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

You can download IBM eServer xSeries, IBM xSeries, and IBM System x publications from the following Web site:

<http://www-304.ibm.com/jct01004c/systems/support/>

Title	Description	Order number
<i>IBM System Storage Productivity Center Introduction and Planning Guide</i>	This guide introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
<i>IBM System Storage Productivity Center Hardware Installation and Configuration Guide</i>	This guide describes how to install and configure the IBM System Storage Productivity Center hardware.	SC23-8822
<i>IBM System Storage Productivity Center Software Installation and User's Guide</i>	This guide describes how to install and use the IBM System Storage Productivity Center software.	SC23-8823

Title	Description	Order number
<i>IBM System Storage Multipath Subsystem Device Driver: User's Guide</i>	This guide describes the IBM System Storage Multipath Subsystem Device Driver Version 1.6 for TotalStorage Products and how to use it with the SAN Volume Controller. This publication is referred to as the <i>IBM System Storage Multipath Subsystem Device Driver: User's Guide</i> .	GC27-2164
<i>IBM TotalStorage DS4300 Fibre Channel Storage Subsystem Installation, User's, and Maintenance Guide</i>	This guide describes how to install and configure the IBM TotalStorage DS4300 Fibre-Channel Storage Subsystem.	GC26-7722
<i>IBM eServer xSeries 306m (Types 8849 and 8491) Installation Guide</i>	This guide describes how to install the IBM eServer xSeries 306m, which is the hardware delivered for some versions of the hardware master console.	MIGR-61615
<i>IBM xSeries 306m (Types 8849 and 8491) User's Guide</i>	This guide describes how to use the IBM eServer xSeries 306m, which is the hardware delivered for some versions of the hardware master console.	MIGR-61901
<i>IBM xSeries 306m (Types 8849 and 8491) Problem Determination and Service Guide</i>	This guide can help you troubleshoot and resolve problems with the IBM eServer xSeries 306m, which is the hardware delivered for some versions of the hardware master console.	MIGR-62594
<i>IBM eServer xSeries 306 (Type 8836) Installation Guide</i>	This guide describes how to install the IBM eServer xSeries 306, which is the hardware delivered for some versions of the hardware master console.	MIGR-55080
<i>IBM eServer xSeries 306 (Type 8836) User's Guide</i>	This guide describes how to use the IBM eServer xSeries 306, which is the hardware delivered for some versions of the hardware master console.	MIGR-55079
<i>IBM eServer xSeries 306 (Types 1878, 8489 and 8836) Hardware Maintenance Manual and Troubleshooting Guide</i>	This guide can help you troubleshoot problems and maintain the IBM eServer xSeries 306, which is the hardware delivered for some versions of the hardware master console.	MIGR-54820

Title	Description	Order number
<i>IBM eServer xSeries 305 (Type 8673) Installation Guide</i>	This guide describes how to install the IBM eServer xSeries 305, which is the hardware delivered for some versions of the hardware master console.	MIGR-44200
<i>IBM eServer xSeries 305 (Type 8673) User's Guide</i>	This guide describes how to use the IBM eServer xSeries 305, which is the hardware delivered for some versions of the hardware master console.	MIGR-44199
<i>IBM eServer xSeries 305 (Type 8673) Hardware Maintenance Manual and Troubleshooting Guide</i>	This guide can help you troubleshoot problems and maintain the IBM eServer xSeries 305, which is the hardware delivered for some versions of the hardware master console.	MIGR-44094
<i>IBM TotalStorage 3534 Model F08 SAN Fibre Channel Switch User's Guide</i>	This guide introduces the IBM TotalStorage SAN Switch 3534 Model F08.	GC26-7454
<i>IBM System x3250 (Types 4364 and 4365) Installation Guide</i>	This guide describes how to install the IBM System x3250, which is the hardware delivered for some versions of the hardware master console.	MIGR-5069761
<i>IBM System x3250 (Types 4364 and 4365) User's Guide</i>	This guide describes how to use the IBM System x3250, which is the hardware delivered for some versions of the hardware master console.	MIGR-66373
<i>IBM System x3250 (Types 4364 and 4365) Problem Determination and Service Guide</i>	This guide can help you troubleshoot and resolve problems with the IBM System x3250, which is the hardware delivered for some versions of the hardware master console.	MIGR-66374
<i>IBM TotalStorage SAN Switch 2109 Model F16 User's Guide</i>	This guide introduces the IBM TotalStorage SAN Switch 2109 Model F16.	GC26-7439
<i>IBM TotalStorage SAN Switch 2109 Model F32 User's Guide</i>	This guide introduces the IBM TotalStorage SAN Switch 2109 Model F32. It also describes the features of the switch and tells you where to find more information about those features.	GC26-7517

Some related publications are available from the following SAN Volume Controller support Web site:

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

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.ibm.com/storage/support/2145
Technical support for IBM storage products	http://www.ibm.com/storage/support/

How to order IBM publications

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

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:

<http://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. SAN Volume Controller overview

The SAN Volume Controller combines software and hardware into a comprehensive, modular appliance that uses symmetric virtualization.

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage subsystems. Those storage systems are then mapped to a set of virtual disks (VDisks) for use by attached host systems. System administrators can view and access a common pool of storage on the SAN. This lets the administrators use storage resources more efficiently and provides a common base for advanced functions.

A SAN is a high-speed fibre-channel network that connects host systems and storage devices. It allows a host system to be connected to a storage device across the network. The connections are made through units such as routers, gateways, hubs, and switches. The area of the network that contains these units is known as the *fabric* of the network.

SAN Volume Controller software

The SAN Volume Controller software performs the following functions for the host systems that attach to SAN Volume Controller over the SAN:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes

The SAN Volume Controller also provides the following functions:

- Large scalable cache
- Copy Services
 - FlashCopy[®] (point-in-time copy)
 - Metro Mirror (synchronous copy)
 - Global Mirror (asynchronous copy)
 - Data migration
- Space management
 - Mapping that is based on desired performance characteristics
 - Metering of service quality
 - Space-efficient logical volumes (thin provisioning)

SAN Volume Controller hardware

Each SAN Volume Controller node is an individual server in a SAN Volume Controller cluster on which the SAN Volume Controller software runs.

The nodes are always installed in pairs, with a minimum of one and a maximum of four pairs of nodes constituting a *cluster*. 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.

I/O groups take the storage that is presented to the SAN by the storage subsystems as MDisks and translates the storage into logical disks, known as VDIsks, that are used by applications on the hosts. A node resides in only one I/O group and provides access to the VDIsks in that I/O group.

The SAN Volume Controller 2145-8G4 is the most current model that is available. In addition, the following models of SAN Volume Controller nodes have been available in previous releases and are still supported with the latest SAN Volume Controller software:

- SAN Volume Controller 2145-8F4
- SAN Volume Controller 2145-8F2
- SAN Volume Controller 2145-4F2

SAN Volume Controller objects

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

Each SAN Volume Controller is a single processing unit called a *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* and each node can be in only one I/O group.

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

Each node does not contain any internal battery backup units and therefore must be connected to an *uninterruptible power supply*, which provides data integrity in the event of a cluster wide power failure. In such situations, the uninterruptible power supply 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 that is presented by backend *disk controllers* as a number of disks, known as *managed disks (MDisks)*.

Each MDisk is divided into a number of *extents* which are numbered, from 0, sequentially from the start to the end of the MDisk. MDisks are collected into groups, known as MDisk groups.

Each VDisk is made up of one or two VDisk copies. Each VDisk copy is an independent physical copy of the data that is stored on the VDisk. A VDisk with two copies is known as a *mirrored VDisk*. VDisk copies are made out of MDisk extents. All the MDisks that contribute to a particular VDisk copy must belong to the same MDisk group.

A VDisk can be space-efficient. This means that the capacity of the VDisk as seen by host systems, called the virtual capacity, can be different from the amount of storage that is allocated to the VDisk from MDisks, called the real capacity. Space-efficient VDIsks can be configured to automatically expand their real capacity by allocating new extents.

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

The nodes detect 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. You can 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 VDisks that have been allocated to them. VDisks can be mapped to a host object. Mapping a VDisk to a host object makes the VDisk accessible to the WWPNs in that host object, and hence the application server itself.

The cluster provides block-level aggregation and volume management for disk storage within the SAN. The cluster manages a number of backend 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 backend physical storage. This prevents any possible conflict between the cluster and the application servers both trying to manage the backend storage.

SAN Volume Controller cluster high availability

A SAN Volume Controller cluster has several features that can be used to deploy a high availability storage system with no single point of failure.

Each I/O group within a cluster consists of a pair of nodes. If a node fails within an I/O group, the other node in the I/O group assumes the I/O responsibilities of the failed node.

If a cluster of SAN Volume Controller nodes is split into two partitions (for example due to a SAN fabric fault), the partition with the majority of nodes continues to process I/O operations. If a cluster is split into two equal-sized partitions, a quorum disk is accessed to determine which half of the cluster continues to read and write data.

Each SAN Volume Controller node has four fibre-channel ports, which allows the node to be attached to multiple SAN fabrics. For high availability, attach the nodes in a cluster to at least two fabrics. SAN Volume Controller software incorporates multipathing software that is used for communication among SAN Volume Controller nodes and for I/O operations among SAN Volume Controller nodes and storage subsystems. If a SAN fabric fault disrupts communication or I/O operations, the multipathing software recovers and retries the operation through an alternative communication path. Also for high availability, configure your host systems to use multipathing software. Then if a SAN fabric fault or node failure occurs, I/O operations among host systems and SAN Volume Controller nodes are retried.

The SAN Volume Controller Virtual Disk Mirroring feature can be used to mirror data across storage subsystems. This feature provides protection against a storage subsystem failure. Although VDisk Mirroring provides additional protection against disk failures, it is not intended to be used as a substitute for RAID in storage subsystems.

The SAN Volume Controller Metro Mirror and Global Mirror features can be used to mirror data between two clusters at different physical locations for disaster recovery.

For very short distances, it is possible to split a cluster between two locations and to use VDisk Mirroring to mirror the data. However, there are configuration restrictions on how a cluster can be split. If one half of the cluster fails, performance is likely to be substantially reduced.

You must configure split clusters so that the following conditions are met:

- Avoid having any interswitch links (ISLs) in the paths between SAN Volume Controller nodes and storage controllers. If it is necessary to have ISLs between SAN Volume Controller nodes and storage controllers, the ISLs should not be oversubscribed as there will be substantial fibre-channel traffic across the ISLs. For most configurations, trunking is required. Because ISL problems are difficult to diagnose, switch-port error statistics must be collected and regularly monitored to detect failures.
- Avoid having any ISLs in the paths between SAN Volume Controller nodes in the same cluster. If it is necessary to have ISLs between SAN Volume Controller nodes in the same cluster, follow these guidelines:
 - At least some ports on each node in the same I/O group must be connected to the same switch in each redundant fabric that is used. Connecting to different blades in a director-class switch is permitted. SAN Volume Controller node-to-node communication in the same I/O group across ISLs is not supported.
 - There should be no more than one ISL hop between SAN Volume Controller nodes in different I/O groups.
 - The ISLs should not be oversubscribed as there will be substantial fibre-channel traffic across the ISLs. For most configurations, trunking is required. Because ISL problems are difficult to diagnose, switch-port error statistics must be collected and regularly monitored to detect failures.
- Use switch zoning to prevent node-to-node traffic using an ISL hop where there are multiple paths across a fabric between SAN Volume Controller nodes, some of which involve an ISL and some of which do not.
- It is best to locate all SAN Volume Controller nodes that are in a cluster in the same rack or adjacent racks. If it is necessary to separate the SAN Volume Controller nodes, follow these rules:
 - All SAN Volume Controller nodes in the same cluster must be connected to the same Ethernet subnet.
 - Some service actions require access to all SAN Volume Controller nodes in a cluster. Therefore, SAN Volume Controller nodes in the same cluster must be reasonably accessible and physically not more than 100 meters apart. Any exceptions must be requested by contacting your IBM representative.
 - A node must be placed in the same rack as the SAN Volume Controller uninterruptible power supply that supplies power to it.

SAN Volume Controller operating environment

To use the SAN Volume Controller, you must meet the minimum hardware and software requirements and ensure that other operating environment criteria are met.

Minimum requirements

You must set up your SAN Volume Controller operating environment according to the following requirements:

- Minimum of one pair of SAN Volume Controller nodes
- Minimum of two uninterruptible power supply units
- One IBM System Storage Productivity Center or one master console per SAN installation for configuration

Features of a SAN Volume Controller 2145-8G4 node

The SAN Volume Controller 2145-8G4 node has the following features:

- 19-inch rack mounted enclosure
- One 4-port 4 Gbps fibre-channel adapter (four fibre-channel ports)
- 8 GB cache memory
- Two dual-core processors

Supported hosts

In a SAN environment, hosts are the application servers that access their data from the storage controllers that are connected to the SAN. Hosts that are running in a number operating environments can connect to the storage through SAN Volume Controller. For a list of the supported operating systems on a host, go to the following Web site:

<http://www.ibm.com/storage/svc>

1. In the **Learn more** column, click **Interoperability**.
2. Click **Recommended software levels** for your SAN Volume Controller code version.
3. Click **Multipathing / Host Drivers, Clustering and SAN Boot Support - By Host Operating System** to view a list of supported operating systems and to access host attachment scripts.

Multipathing software

For the most current information, go to the following Web site:

<http://www.ibm.com/storage/svc>

1. In the **Learn more** column, click **Interoperability**.
2. Click **Recommended software levels** for your SAN Volume Controller code version.
3. Click **Multipathing / Host Drivers, Clustering and SAN Boot Support - By Host Operating System** to view a list of supported operating systems and to access multipath drivers. You can also view **Multipath Driver Co-existence with SDD** information.

User interfaces

The SAN Volume Controller software provides the following user interfaces:

- The 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) that uses Secure Shell (SSH)

Application programming interfaces

The SAN Volume Controller software provides an application programming interface called the Common Information Model (CIM) agent, which supports the Storage Management Initiative Specification (SMI-S) of the Storage Network Industry Association.

2145-1U uninterruptible power supply

A 2145-1U uninterruptible power supply is used exclusively to maintain data that is held in the SAN Volume Controller dynamic random access memory (DRAM) in the event of an unexpected loss of external power. This use differs from the traditional uninterruptible power supply that enables continued operation of the device that it supplies when power is lost.

With a 2145-1U uninterruptible power supply, data is saved to the internal disk of the SAN Volume Controller node. The uninterruptible power supply units are required to power the SAN Volume Controller nodes even if the input power source is considered uninterruptible.

The uninterruptible power supply maintains continuous SAN Volume Controller-specific communications with its attached SAN Volume Controller nodes. A SAN Volume Controller node cannot operate without the uninterruptible power supply. The uninterruptible power supply must be used in accordance with documented guidelines and procedures and must not power any equipment other than a SAN Volume Controller node. Each uninterruptible power supply must be in the same rack as the node it powers.

Redundant ac power switch

The redundant ac power switch is an optional feature that makes the SAN Volume Controller nodes more resilient to power failure. The redundant ac power switch is not a replacement for an uninterruptible power supply. You must still use a 2145-1U uninterruptible power supply for each node.

You must connect the redundant ac power switch to two independent power circuits. One power circuit connects to the main power input port and the other power circuit connects to the backup power input port. If the main power to the SAN Volume Controller node fails for any reason, the redundant ac power switch automatically uses the backup power source. When power is restored, the redundant ac power switch automatically changes back to using the main power source.

Place the redundant ac power switch in the same rack as the SAN Volume Controller node. The redundant ac power switch logically sits between the rack power distribution unit and the 2145-1U uninterruptible power supply.

You can use a single redundant ac power switch to power one or two SAN Volume Controller nodes. If you use the redundant ac power switch to power two nodes, the nodes must be in different I/O groups. In the event that the redundant ac power switch fails or requires maintenance, both nodes power off. Because the nodes are in two different I/O groups, the hosts do not lose access to the backend disk data.

For maximum resilience to failure, use one redundant ac power switch to power each SAN Volume Controller node.

The following graphic shows a redundant ac power switch.



IBM System Storage Productivity Center

The IBM System Storage Productivity Center (SSPC) is an integrated hardware and software solution that provides a single point of entry for managing SAN Volume Controller clusters, IBM System Storage DS8000 systems, and other components of your data storage infrastructure.

SSPC simplifies storage management in the following ways:

- Centralizing the management of storage network resources with IBM storage management software
- Providing greater synergy between storage management software and IBM storage devices
- Reducing the number of servers that are required to manage your software infrastructure
- Providing simple migration from basic device management to storage management applications that provide higher-level functions

SSPC includes the following software components:

- SAN Volume Controller Console, including the CIM agent
- PuTTY (SSH client software)
- IBM TotalStorage Productivity Center Basic Edition, which can be used to access the IBM System Storage DS8000 Storage Manager
- DB2® Enterprise Server Edition

Figure 1 on page 8 shows an overview of how SSPC and the components of IBM TotalStorage Productivity Center, IBM System Storage DS8000, and SAN Volume Controller interrelate with each other.

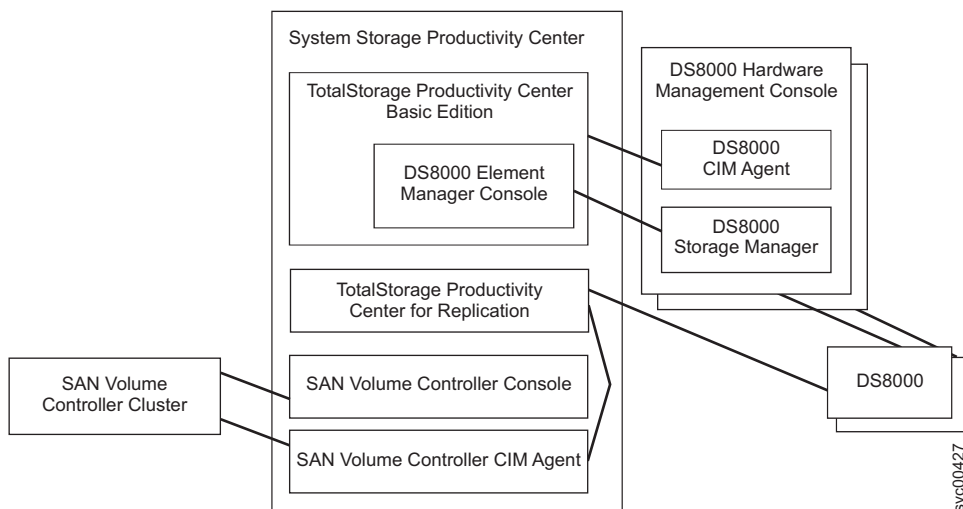


Figure 1. Overview of the IBM System Storage Productivity Center

For more information on SSPC, see the *IBM System Storage Productivity Center Introduction and Planning Guide*.

Assist On-site and remote service

When you contact IBM to help you resolve a problem with your SAN Volume Controller environment, the IBM service representative might suggest using the IBM Assist On-site tool to remotely access the IBM System Storage Productivity Center (SSPC) or master console. This type of remote service can help you reduce service costs and shorten repair times.

The IBM Assist On-site tool is a remote desktop-sharing solution that is offered through the IBM Web site. With it, the IBM service representative can remotely view your system to troubleshoot a problem. You can maintain a chat session with the IBM service representative so that you can monitor the activity and either understand how to fix the problem yourself or allow the representative to fix it for you.

To use the IBM Assist On-site tool, the SSPC or master console must be able to access the Internet. The following Web site provides further information about this tool:

<http://www.ibm.com/support/assistsite/>

When you access the Web site, you sign in and enter a code that the IBM service representative provides to you. This code is unique to each IBM Assist On-site session. A plug-in is downloaded onto your SSPC or master console to connect you and your IBM service representative to the remote service session. The IBM Assist On-site contains several layers of security to protect your applications and your computers. You can also use security features to restrict access by the IBM service representative.

Your IBM service representative can provide you with more detailed instructions for using the tool.

Secure Shell protocol through PuTTY

Secure Shell (SSH) software is a client-server protocol that can be used from the IBM System Storage Productivity Center or from a host server to enable you to control the SAN Volume Controller through a command-line interface (CLI).

SSH provides a secure communications channel between systems. You can configure SSH to use a key pair (a private key and a public key) to establish the secure connection to a remote system. If you want to create an SSH connection (such as to the SAN Volume Controller cluster), you must place the public key on every system. The SAN Volume Controller provides a command to distribute a key to each node in the cluster.

Data and event notifications

The SAN Volume Controller can use SNMP traps, Call Home e-mail, and Inventory Information e-mail to provide necessary data and event notifications to you and to the IBM Support Center.

The following types of information are sent from the SAN Volume Controller:

- Simple Network Management Protocol (SNMP) traps
- Call Home e-mail
- Inventory information

Simple Network Management Protocol traps

Simple network management protocol (SNMP) is the standard protocol for managing networks and exchanging messages. SNMP enables the SAN Volume Controller to send external messages that notify personnel about an event. An SNMP manager allows you to view the messages that the SNMP agent sends. You can use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure and modify your SNMP settings. SNMP traps and Call Home e-mail can be sent simultaneously.

Call Home e-mail

The Call Home feature allows the transmission of operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification e-mail. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues.

You must configure an SMTP server to be able to send e-mail outside of your local area network. The SMTP server must allow the relaying of e-mail from the SAN Volume Controller cluster IP address. You can then use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure the e-mail settings, including contact information and e-mail recipients. For compatibility with other SMTP servers, ensure that you set the reply address to a valid e-mail address. Send a test e-mail to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time using the SAN Volume Controller Console or the SAN Volume Controller command-line interface.

Inventory information e-mail

Inventory information e-mail is a type of Call Home notification. Inventory information can be sent to IBM to assist IBM service personnel in evaluating your SAN Volume Controller system. Because inventory information is sent using the Call Home e-mail function, you must meet the Call Home function requirements and enable the Call Home e-mail function before you can attempt to send inventory information e-mail. You can adjust the contact information, adjust the frequency of inventory e-mail, or manually send an inventory e-mail using the SAN Volume Controller Console or the SAN Volume Controller command-line interface. Inventory information is automatically reported to IBM when you activate error reporting.

Clusters

All of your configuration and service tasks 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.

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.

All configurations are replicated across all nodes in the cluster; however, only some service actions can be performed at the node level. Because configuration is performed at the cluster level, an IP address is assigned to the cluster instead of each node.

Cluster state

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

The cluster state information is held in nonvolatile memory. If the mainline power fails, the 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 input/output (I/O) group that are using that information. 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 2 on page 11 shows an example of a cluster that contains 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 (1), that results in a change being made to the configuration. The configuration node controls updates to the cluster (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 (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 2. Cluster, nodes, and cluster state.

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. The configuration 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 input/output (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 virtual disks (VDisks) that are managed by the I/O group.

Note: The cluster can continue to run without loss of access to data as long as one node from each I/O group is available.

A tiebreak 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 tiebreak exists.

The cluster automatically chooses three managed disks (MDisks) to be candidate *quorum disks* and assigns them quorum indexes of 0, 1, and 2. One of these disks is used to settle a tiebreak condition.

If a tiebreak 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 disk at any time by issuing the following command:

```
svctask setquorum
```

I/O groups and uninterruptible power supply units

Each pair of nodes is known as an *input/output (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 back-end storage and all application servers are visible to all of the I/O groups. Each pair of nodes has the responsibility to serve I/O operations on a particular virtual disk (Vdisk).

VDisks are logical disks that are presented to the SAN by SAN Volume Controller nodes. VDisks 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 VDisk, it can access the VDisk with either of the nodes in the I/O group. When you create a VDisk, you can specify a preferred node. If you do not specify a preferred node, one is automatically assigned after the VDisk is created. The preferred node is used to access the VDisk.

Each I/O group only has two nodes. The distributed cache inside the SAN Volume Controller is replicated across both nodes in the I/O group. When I/O is performed to a VDisk, 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 VDisk 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 you can easily scale the I/O capability of SAN Volume Controller clusters by adding I/O groups to obtain additional throughput.

Figure 3 shows an example I/O group. A write operation from a host is shown (item 1), that is targeted for VDisk A. This write operation is targeted at the preferred node, Node 1 (item 2). The write operation is cached and a copy of the data is made in the partner node, the Node 2 cache (item 3). The write operation is now complete so far as the host is concerned. At some time later the data is written, or de-staged, 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 3. 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 write operations for the VDisk 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 VDisk that are assigned to the I/O group cannot be accessed.

When a VDisk is created, you must specify the I/O group that will provide access to the VDisk. However, VDIsks can be created and added to I/O groups that contain offline nodes. I/O access is not 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 VDIsks to the recovery I/O group and then into a working I/O group. I/O access is not possible when VDIsks are assigned to the recovery I/O group.

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 you are 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 that is filtered by controller. This view enables you to associate the MDisks with the RAID 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 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 can be more than one controller), can be used to associate the MDisks in the cluster with the RAID presented by the controllers.

To minimize data loss, virtualize only those RAID 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 VDIsks.

Virtual Disk Mirroring feature

The VDisk Mirroring feature allows a VDisk to have two physical copies. Each VDisk copy can belong to a different MDisk group. Each copy has the same virtual capacity as the VDisk.

When a server writes to a mirrored VDisk, the SAN Volume Controller cluster writes the data to both copies. When a server reads a mirrored VDisk, the SAN Volume Controller cluster picks one of the copies to read. If one of the mirrored VDisk copies is temporarily unavailable; for example, because the RAID controller that provides the MDisk group is unavailable, the VDisk remains accessible to servers. The SAN Volume Controller cluster remembers which areas of the VDisk are written and resynchronizes these areas when both copies are available.

You can create a VDisk with one or two copies and convert a non-mirrored VDisk into a mirrored VDisk by adding a copy. When a copy is added in this way, the SAN Volume Controller cluster synchronizes the new copy so that it is the same as the existing VDisk. Servers can access the VDisk during this synchronization process.

You can convert a mirrored VDisk into a non-mirrored VDisk by deleting one copy or by splitting one copy to create a new non-mirrored VDisk.

The VDisk copy can be any type: image, striped, sequential, and space-efficient or not. The two copies can be of completely different types.

VDisk Mirroring can be used for the following applications:

- Improve availability of VDIs by protecting them from a single storage controller failure.
- Allow concurrent maintenance of a storage controller that does not natively support concurrent maintenance.
- Provide an alternative method of data migration with better availability characteristics. While a VDisk is being migrated using the data migration feature, it is vulnerable to failures on both the source and target MDisk group. VDisk Mirroring provides an alternative because you can start with a non-mirrored VDisk in the source MDisk group and then add a copy to that VDisk in the destination MDisk group. When the VDisk is synchronized, you can delete the original copy that is in the source MDisk group. During the synchronization process, the VDisk remains available even if there is a problem with the destination MDisk group.

When you use VDisk Mirroring, consider how quorum candidate disks are allocated. VDisk Mirroring maintains some state data on the quorum disks. If the quorum disks are not accessible and VDisk Mirroring is unable to update the state information, a mirrored VDisk might need to be taken offline to maintain data integrity. To ensure the high availability of the system, ensure that multiple quorum candidate disks, allocated on different controllers, are configured.

Space-Efficient Virtual Disk feature

When you create a virtual disk (VDisk), you can designate it as space-efficient. A space-efficient VDisk has a virtual capacity and a real capacity.

The VDisk consumes physical storage that is equivalent to the real capacity but appears to host systems to have a different capacity, which is the virtual capacity.

Typically the virtual capacity is significantly larger than the real capacity. The SAN Volume Controller cluster uses the real capacity to store data that is written to the virtual capacity and additional information that describes the parts of the virtual capacity that have been written. As more information is written to the virtual capacity, more of the real capacity is used. The SAN Volume Controller cluster identifies reads to unwritten parts of the virtual capacity and returns zeroes to the server without using any of the real capacity.

The SAN Volume Controller must maintain extra metadata that describes the contents of space-efficient VDIs. This means the I/O rates that are obtained from space-efficient VDIs are slower than those obtained from fully allocated VDIs that are allocated on the same MDIs.

Space-efficient VDIs can also simplify server administration. Instead of assigning a VDI with some capacity to an application and increasing that capacity as the application's needs change, you can configure a VDI with a large virtual capacity for the application and then increase or shrink the real capacity as the application needs change, without disrupting the application or server.

You can specify to autoexpand space-efficient VDIs. The real capacity is then expanded as the capacity is used. During this process, a fixed amount of unused real capacity is maintained. This amount is known as the *contingency capacity*. When you create a space-efficient VDI, the contingency capacity is initially set to the same amount as the real capacity. If you manually change the real capacity, the contingency capacity becomes the difference between the used capacity and the new real capacity.

A space-efficient VDI can be configured to generate a warning event when the used real capacity exceeds a specified amount or percentage of the virtual capacity. You can use the warning event to trigger other actions, such as taking low-priority applications offline or migrating data into other MDI groups.

If a space-efficient VDI does not have enough real capacity for a write operation, the VDI is taken offline and an error is logged (error ID 060001). Access to the VDI can be restored by increasing its real capacity. Then, when sufficient real capacity is available, the error is automatically marked as fixed and access is restored. However, if that VDI is configured to automatically expand its real capacity, you must mark the error as fixed before access is restored.

Data migration

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

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

Applications for data migration

The following lists the different applications for data migration:

- Redistribution of workload within a cluster across MDIs. You can redistribute the workload, by one of the following ways:
 - 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.

Cluster configuration 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 backup the cluster configuration settings to tertiary storage. You can use the configuration backup functions to backup the cluster configuration.

For complete disaster recovery, regularly backup the business data that is stored on virtual disks at the application server level or the host level.

Copy Services features

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

The following Copy Services features 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 that the copy is continuously updated.

Global Mirror

Provides a consistent copy of a source VDisk on a target VDisk. Data is written to the target VDisk asynchronously, so that the copy is continuously updated, but the copy might not contain the last few updates in the event that a disaster recovery operation is performed.

FlashCopy

FlashCopy is a Copy Services feature that is available with the SAN Volume Controller.

In its basic mode, the FlashCopy feature copies the contents of a source virtual disk (VDisk) to a target VDisk. Any data that existed on the target VDisk is lost and is replaced by the copied data. After the copy operation has completed, the

target VDIsks contain the contents of the source VDIsks as they existed at a single point in time unless target writes have been performed. The FlashCopy feature 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 to complete, the resulting data on the target VDisk is presented so that the copy appears to have occurred immediately.

Although it is difficult to make a consistent copy of a data set that is constantly updated, point-in-time copy techniques 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, the resulting copy might contain data that 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 copied, the copy contains the referenced object at its new location but the copied reference still points to the old location.

More advanced FlashCopy features allow operations to occur on multiple source and target VDIsks. FlashCopy management operations are coordinated to allow a common single point in time for copying target VDIsks from their respective source VDIsks. This allows a consistent copy of data that spans multiple VDIsks. The FlashCopy feature also allows multiple target VDIsks to be copied from each source VDisk. This can be used to create images from different points in time for each source VDisk.

The Cascaded FlashCopy feature also allows a FlashCopy target VDisk to be the source VDisk of another FlashCopy mapping. You also have the option of using the Incremental FlashCopy feature, which potentially reduces the amount of time to complete the copy operation after the initial copy has completed. Only the differences are copied when the FlashCopy mapping restarts.

FlashCopy associates a source VDisk and a target VDisk in a FlashCopy Mapping. The source VDIsks and target VDIsks must meet the following requirements:

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

Any VDisk that is part of a FlashCopy operation can be space-efficient. Using a space-efficient FlashCopy target can reduce the amount of storage that is required to maintain a point-in-time copy. The source VDIsks and target VDIsks can also be mirrored to improve availability of the VDIsks.

Metro Mirror and Global Mirror

The Metro Mirror and Global Mirror Copy Services features enable you to set up a relationship between two virtual disks (VDIsks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

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

For both Metro Mirror and Global Mirror copy types, 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 I/O operations to the secondary VDisk.

The Metro Mirror feature provides a *synchronous*-copy process. When a host writes to the primary VDisk, it does not receive confirmation of I/O completion until the write operation has completed for the copy on both the primary VDisk and the secondary VDisk. This ensures that the secondary VDisk is always up-to-date with the primary VDisk in the event that a failover operation must be performed. However, the host is limited to the latency and bandwidth limitations of the communication link to the secondary VDisk.

The Global Mirror feature provides an *asynchronous*-copy process. When a host writes to the primary VDisk, confirmation of I/O completion is received before the write operation has completed for the copy on the secondary VDisk. If a failover operation is performed, the application must recover and apply any updates that were not committed to the secondary VDisk. If I/O operations on the primary VDisk are paused for a small length of time, the secondary VDisk can become an exact match of the primary VDisk.

The Metro Mirror and Global Mirror operations support the following functions:

- 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 and Global Mirror relationships with itself and one other cluster.

- Intercluster and intracluster Metro Mirror and Global Mirror relationships can be used concurrently within a cluster.
- The intercluster link is bidirectional. This means that it can copy data from cluster A to cluster B for one pair of VDIsks while copying data from cluster B to cluster A for a different pair of VDIsks.
- The copy direction can be reversed for a consistent relationship.
- Consistency groups are supported to manage a group of relationships that must be kept synchronized for the same application. This also simplifies administration, because a single command that is issued to the consistency group is applied to all the relationships in that group.

Chapter 2. SAN Volume Controller physical installation planning

Before the IBM service representative can set up your SAN Volume Controller environment, you must verify that the prerequisite conditions for the SAN Volume Controller, the uninterruptible power supply, and optional redundant ac power switch installation are met.

This information considers only the models that are currently available to purchase. For information on all supported models, see the *IBM System Storage SAN Volume Controller: Service Guide*.

1. Does your physical site meet the environment requirements for the SAN Volume Controller, uninterruptible power supply, and the redundant ac power switch?
2. Do you have adequate rack space for your hardware? Ensure you have the following rack space for your components:
 - The SAN Volume Controller: One Electronic Industries Alliance (EIA) unit high for each node.
 - The 2145-1U uninterruptible power supply: One EIA unit high for each 2145-1U uninterruptible power supply.
 - If you are using a redundant ac power switch: One EIA unit high for each redundant ac power switch
3. Do the power circuits you will be using have sufficient capacity and the correct sockets for your SAN Volume Controller installation?

A clearly visible and accessible emergency power off switch is required.

If you are using the redundant ac power switch, two independent power circuits are required. One circuit connects to the redundant ac power switch main input and the other circuit connects to the redundant ac power switch backup input.
4. Have you provided appropriate connectivity by preparing your environment?
5. Do you have a keyboard and display available in the event that a service action is required? The SAN Volume Controller 2145-8G4 requires a USB keyboard.
6. Do you have adequate space available and power for the IBM System Storage Productivity Center?

SAN Volume Controller 2145-8G4 environment requirements

Before the SAN Volume Controller 2145-8G4 is installed, the physical environment must meet certain requirements. This includes verifying that adequate space is available and that requirements for power and environmental conditions are met.

Input-voltage requirements

Ensure that your environment meets the following voltage requirements.

Voltage	Frequency
200 to 240 V single phase ac	50 or 60 Hz

Power requirements for each node

Ensure that your environment meets the following power requirements.

The power that is required depends on the node type, the uninterruptible power supply type, and whether the redundant ac power feature is used.

Components	Power requirements
SAN Volume Controller 2145-8G4 + 2145-1U uninterruptible power supply	470 W

For each redundant ac power switch, add 20 W to the power requirements.

Circuit breaker requirements

The 2145-1U uninterruptible power supply has an integrated circuit breaker and does not require additional protection.

Environment requirements without redundant ac power

Ensure that your environment falls within the following ranges if you are not using redundant ac power.

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 (73°F)
Operating in higher altitudes	10°C to 32°C (50°F to 90°F)	914 to 2133 m (2998 to 6988 ft)	8% to 80% noncondensing	23°C (73°F)
Powered off	10°C to 43°C (50°F to 110°F)	0 to 2133 m (2998 to 6988 ft)	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)

Environment requirements with redundant ac power

Ensure that your environment falls within the following ranges if you are using redundant ac power.

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	15°C to 32°C (59°F to 90°F)	0 to 914 m (0 to 2998 ft)	20% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	15°C to 32°C (59°F to 90°F)	914 to 2133 m (2998 to 6988 ft)	20% to 80% noncondensing	23°C (73°F)
Powered off	10°C to 43°C (50°F to 110°F)	0 to 2133m (0 to 6988 ft)	20% to 80% noncondensing	27°C (81°F)

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
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)

Preparing your environment

The following tables list the physical characteristics of the SAN Volume Controller 2145-8G4 node.

Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the node.

Height	Width	Depth	Maximum weight
43 mm (1.69 in.)	440 mm (17.32 in.)	686 mm (27 in.)	12.7 kg (28 lb)

Additional space requirements

Ensure that space is also available in the rack for the following additional space requirements around the node.

Location	Additional space requirements	Reason
Left and right sides	50 mm (2 in.)	Cooling air flow
Back	Minimum: 100 mm (4 in.)	Cable exit

Heat output of each SAN Volume Controller 2145-8G4 node

The node dissipates the following maximum heat output.

Model	Heat output per node
SAN Volume Controller 2145-8G4	400 W (1350 Btu per hour)

Requirements for the uninterruptible power supply environment

An uninterruptible power supply environment requires that certain specifications for the physical site of the SAN Volume Controller must be met.

2145-1U uninterruptible power supply environment

The SAN Volume Controller 2145-8G4, SAN Volume Controller 2145-8F4, SAN Volume Controller 2145-8F2, and SAN Volume Controller 2145-4F2 are supported with the 2145-1U uninterruptible power supply.

2145-1U uninterruptible power supply specifications

The following tables describe the physical characteristics of the 2145-1U uninterruptible power supply.

2145-1U uninterruptible power supply dimensions and weight

Ensure that space is available in a rack that is capable of supporting the 2145-1U uninterruptible power supply.

Height	Width	Depth	Maximum weight
44 mm (1.73 in.)	439 mm (17.3 in.)	579 mm (22.8 in.)	16 kg (35.3 lb)
Note: The 2145-1U uninterruptible power supply package, which includes support rails, weighs 18.8 kg (41.4 lb).			

Heat output

The 2145-1U uninterruptible power supply unit produces the following approximate heat output.

Model	Heat output during normal operation	Heat output during battery operation
2145-1U uninterruptible power supply	40 W (135 Btu per hour)	150 W (510 Btu per hour)

Power cables for the 2145-1U uninterruptible power supply

If you do not connect the 2145-1U uninterruptible power supply to a rack power distribution unit (PDU) or redundant ac power switch, you must follow your country or region's power requirements to choose the appropriate power cable for the 2145-1U uninterruptible power supply.

The 2145-1U uninterruptible power supply is supplied with an IEC 320-C13 to C14 jumper to connect it to a rack PDU. You can also use this cable to connect the 2145-1U uninterruptible power supply to the redundant ac power switch.

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

Country or region	Length	Connection type (attached plug designed for 200 - 240 V ac input)	Part number
United States of America (Chicago), Canada, Mexico	1.8 m (6 ft)	NEMA L6-15P	39M5115

Country or region	Length	Connection type (attached plug designed for 200 - 240 V ac input)	Part number
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	39M5116
Afghanistan, Algeria, Andorra, Angola, Austria, Belgium, Benin, Bulgaria, Burkina Faso, Burundi, Cameroon, Central African Republic, Chad, Czech Republic, Egypt, Finland, France, French Guiana, Germany, Greece, Guinea, Hungary, Iceland, Indonesia, Iran, Ivory Coast, Jordan, Lebanon, Luxembourg, Macao S.A.R. of China, 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	39M5123
Antigua, Bahrain, Brunei, Channel Islands, Hong Kong S.A.R. of China, 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	39M5151
Argentina	2.8 m (9 ft)	IRAM 2073	39M5068
Argentina, Australia, New Zealand, Papua New Guinea, Paraguay, Uruguay, Western Samoa	2.8 m (9 ft)	AS/NZS 3112/2000	39M5102
Bangladesh, Burma, Pakistan, South Africa, Sri Lanka	2.8 m (9 ft)	SABS 164	39M5144
Chile, Ethiopia, Italy, Libya, Somalia	2.8 m (9 ft)	CEI 23-16	39M5165
People's Republic of China	2.8 m (9 ft)	GB 2099.1	39M5206
Denmark	2.8 m (9 ft)	DK2-5a	39M5130
Israel	2.8 m (9 ft)	SI 32	39M5172

Country or region	Length	Connection type (attached plug designed for 200 - 240 V ac input)	Part number
Liechtenstein, Switzerland	2.8 m (9 ft)	IEC 60884 Stnd. Sheet 416534?2 (CH Type 12)	39M5158
Thailand	2.8 m (9 ft)	NEMA 6-15P	39M5095

Preparing your redundant ac power environment

Ensure that your physical site meets the installation requirements for the redundant ac power switch.

The redundant ac power switch requires two independent power sources that are provided through two rack-mounted power distribution units (PDUs). The PDUs must have IEC320-C13 outlets.

The redundant ac power switch comes with two IEC 320-C19 to C14 power cables to connect to rack PDUs. There are no country-specific cables for the redundant ac power switch.

The power cable between the redundant ac power switch and the 2145-1U uninterruptible power supply is rated at 10 A.

Redundant ac power switch specifications

The following tables list the physical characteristics of the redundant ac power switch.

Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the redundant ac power switch.

Height	Width	Depth	Maximum weight
43 mm (1.69 in.)	192 mm (7.56 in.)	240 mm	2.6 kg (5.72 lb)

Additional space requirements

Ensure that space is also available in the rack for the side mounting plates on either side of the redundant ac power switch.

Location	Width	Reason
Left side	124 mm (4.89 in.)	Side mounting plate
Right side	124 mm (4.89 in.)	Side mounting plate

Heat output (maximum)

The maximum heat output that is dissipated inside the redundant ac power switch is approximately 20 watts (70 Btu per hour).

Connections

Ensure that you are familiar with the specific connection types for the SAN Volume Controller and the uninterruptible power supply.

Each SAN Volume Controller requires the following connections:

- Each SAN Volume Controller node requires one Ethernet cable to connect it to an Ethernet switch or hub. A 10/100/1000 Mb Ethernet connection is required. Both Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) are supported.
- To ensure cluster failover operations, all nodes in a cluster must be connected to the same IP subnet.
- 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) gigabit interface converters (GBICs) for connection to a fibre-channel switch.

Each uninterruptible power supply requires a serial cable that connects the uninterruptible power supply to the SAN Volume Controller node. Ensure that for each node, the serial and power cables come from the same uninterruptible power supply.

Preparing your IBM System Storage Productivity Center environment

Before an IBM service representative can install the IBM System Storage Productivity Center, you must prepare the physical environment. This includes verifying that adequate space is available and that requirements for power and environmental conditions are met.

For information about setting up the physical environment, see the *IBM System Storage Productivity Center Introduction and Planning Guide*.

Web browser requirements to access the SAN Volume Controller Console

Although an IBM System Storage Productivity Center installation includes a Web browser, other browsers are also supported. You must have a supported Web browser to access the SAN Volume Controller Console.

The Web browsers that you can use to access the SAN Volume Controller Console depend on the following conditions:

- Whether you access the SAN Volume Controller Console directly or from a remote system.
- The operating system from which you access the SAN Volume Controller Console.

For a list of the supported browsers, see the *Recommended Software Levels* document at the following Web site:

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

Physical configuration planning of a system with SAN Volume Controller

Before you or your IBM service representative installs the SAN Volume Controller nodes, uninterruptible power supply units, and IBM System Storage Productivity Center, you must plan the physical configuration and the initial settings for the system.

1. Download the hardware location chart, the cable connection table, the configuration data table, and the optional redundant ac power switch connection chart from the following Web site:
<http://www.ibm.com/storage/support/2145>
 - a. Click the **Plan/upgrade** tab.
 - b. Under **Sizing/Capacity**, click **SAN Volume Controller planning**.
 - c. Click your language to view the Planning Guide abstract. The charts and tables are under the **Related information** heading.
2. Use the hardware location chart to record the physical configuration of your system.
3. Use the cable connection table to record how you want your SAN Volume Controller units, the uninterruptible power supply units, and the IBM System Storage Productivity Center to be connected.
4. Use the configuration data table to record the data that you and the IBM service representative need before the initial installation.

When you or your IBM service representative have completed these tasks, you can perform the physical installation.

Requirements and guidelines for completing the hardware location chart

The hardware location chart represents the rack into which the SAN Volume Controller nodes are installed. Each row of the chart represents one Electronic Industries Alliance (EIA) 19-inch rack space.

When you design the rack for the SAN Volume Controller, use the hardware location chart to record the physical configuration of your system.

You can download the hardware location chart from the following Web site:

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

To select the hardware location chart, follow these steps:

1. Click the **Plan/upgrade** tab.
2. Under **Sizing/Capacity**, click **SAN Volume Controller planning**.
3. Click your language to view the Planning Guide abstract.
4. Under **Related information**, click **SAN Volume Controller - Hardware Location Chart**.

When you complete the information in the hardware location chart, consider these requirements:

- Do not exceed the maximum power rating of the rack and input power supply.

- A SAN Volume Controller node is one EIA unit high. Therefore, for each SAN Volume Controller node that is to be installed, fill in the row that represents the position that the SAN Volume Controller node is to occupy.
- The 2145 uninterruptible power supply is two EIA units high. Therefore, for each 2145 uninterruptible power supply, fill in two rows.
- The 2145-1U uninterruptible power supply is one EIA unit high. Therefore, for each 2145-1U uninterruptible power supply, fill in one row.
- The redundant ac power switch is one EIA unit high. Therefore, for each redundant ac power switch, fill in one row.
- The IBM System Storage Productivity Center hardware 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.
- A cluster can contain no more than eight SAN Volume Controller nodes.
- Install the SAN Volume Controller nodes in pairs to provide redundancy and concurrent maintenance.
- The IBM service representative 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 addition to the requirements for completing your information in the hardware location chart, consider also the following guidelines and recommendations.

- Position SAN Volume Controller nodes so that information on the display screen is easily viewed and so that the controls that are used to navigate the display menu are easily reached. Place SAN Volume Controller nodes in EIA 11-36.
- To view the front panels of the SAN Volume Controller node while you perform maintenance procedures, place the IBM System Storage Productivity Center in close proximity to the SAN Volume Controller cluster. To ensure that the IBM System Storage Productivity Center is at a convenient height for viewing the monitor and using the keyboard, place the IBM System Storage Productivity Center in EIA 17-24. Position the IBM System Storage Productivity Center server, the keyboard, and the monitor unit close to each other. Ensure you allow access to the CD drive that is located in the server while the keyboard and monitor unit are pulled out and are open. To allow the monitor to open, do not place the keyboard and monitor unit directly below a SAN Volume Controller node.
- To reduce the chance of a simultaneous input power failure at both uninterruptible power supply units, either use the redundant ac power switch or connect each uninterruptible power supply unit to a separate electrical power source on a separate branch circuit.
- You might want to put the switches between the SAN Volume Controller nodes. The uninterruptible power supply units must remain in the lowest position of the rack.

Requirements for completing the cable connection table

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

You can download the cable connection table from the following Web site:

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

To select the cable connection table, follow these steps:

1. Click the **Plan/upgrade** tab.
2. Under **Sizing/Capacity**, click **SAN Volume Controller planning**.
3. Click your language to view the Planning Guide abstract.
4. Under **Related information**, click **SAN Volume Controller - Cable Connection Table**.

The following terms and descriptions assist you in completing the cable connection table for the SAN Volume Controller nodes:

Term	Description
Uninterruptible power supply	The uninterruptible power supply to which the SAN Volume Controller node is connected.
Ethernet	The Ethernet hub or switch to which the SAN Volume Controller node is connected. The Ethernet hub or switch must have a connection to the IBM System Storage Productivity Center. If you need to generate e-mail reports, there must be a connection to a mail server.
Fibre-channel ports 1 through 4	The fibre channel switch ports to which the four fibre channel ports of the SAN Volume Controller node are connected. When you view it from the back of the SAN Volume Controller node, the ports are numbered 1 through 4, from left to right. Ignore the markings on the back of the SAN Volume Controller node.

For the IBM System Storage Productivity Center hardware, use the following terms and descriptions to complete the cable connection table:

Term	Description
Ethernet to SAN Volume Controller	The Ethernet port used to connect the IBM System Storage Productivity Center and an Ethernet switch or router.
Ethernet to Internet for remote service or upgrade procedures	(Optional) The Ethernet port used to connect the IBM System Storage Productivity Center to the Internet for remote service or upgrade. You can determine how you want to provide the port. For example, you can provide the set up and configure the port by using any of the following methods: <ul style="list-style-type: none">• Provide Internet access through the same port that is used to access the SAN Volume Controller nodes. This method does not provide as much security as the other example methods.• Use an alternate Ethernet port. For added security if you choose this method, you can disconnect this port and connect it only when a remote service connection is needed.• Choose not to provide the IBM System Storage Productivity Center with any access to the public Internet and allow the IBM® Assist On-site tool to access another console that can access the IBM System Storage Productivity Center Web server over an intranet.

Configuration data table guidelines

You must fill out the configuration data table before the SAN Volume Controller nodes and IBM System Storage Productivity Center can be installed.

You can download the configuration data table from the following Web site:

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

To select the configuration data table, follow these steps:

1. Click the **Plan/upgrade** tab.
2. Under **Sizing/Capacity**, click **SAN Volume Controller planning**.
3. Click your language to view the Planning Guide abstract.
4. Under **Related information**, click **SAN Volume Controller - Configuration Data Table**.

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 setting is English.
Cluster IPv4 address or Cluster IPv6 address	The address that is used for all typical configuration and service access to the cluster.
Service IPv4 address or Service IPv6 address	The address that is used for service mode access to a node from the cluster.
Gateway IPv4 address or Gateway IPv6 address	The IP address for the default local gateway for the cluster.
IPv4 subnet mask	The subnet mask of the cluster, which identifies the IPv4 network on which the cluster is operating.
IPv6 prefix	The prefix of the cluster, which identifies the IPv6 network on which the cluster is operating.

Include the following information for the IBM System Storage Productivity Center.

Term	Description
Machine name	The name of the IBM System Storage Productivity Center. This must be a fully qualified DNS name.
IBM System Storage Productivity Center IP addresses	The addresses that are used for access to the IBM System Storage Productivity Center.
IBM System Storage Productivity Center gateway IP address	The IP address for the local gateway for the IBM System Storage Productivity Center.
IBM System Storage Productivity Center subnet mask	The subnet mask for the IBM System Storage Productivity Center.

Requirements for completing the redundant ac power switch connection chart

If you are using the optional redundant ac power switch feature, you must fill out the redundant ac power switch connection chart before this feature can be installed.

You can download the redundant ac power switch connection chart from the following Web site:

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

To select the redundant ac power switch connection chart, follow these steps:

1. Click the **Plan/upgrade** tab.
2. Under **Sizing/Capacity**, click **SAN Volume Controller planning**.
3. Click your language to view the Planning Guide abstract.
4. Under **Related information**, click **SAN Volume Controller - Redundant AC Power Chart**.

Determine if you want to use the redundant ac power switch to power one or two nodes. If you plan to power two nodes, the nodes must be in different I/O groups. Therefore, a cluster with a single I/O group requires two redundant ac power switch units.

You must plan the route of the input power cables for the redundant ac power switch units. These cables are connected to the front edge of the redundant ac power switch and must be routed to the power distribution unit for the rack. You might have to leave a free slot in the rack to allow the cables to run from front to back.

Cabling of redundant ac power switch (example)

You must properly cable the redundant ac power switch units in your environment.

Note: While this topic provides an example of the cable connections, it does not indicate a preferred physical location for the components.

Figure 4 on page 31 shows an example of the main wiring for a SAN Volume Controller cluster with the redundant ac power switch feature. The four-node cluster consists of two I/O groups:

- I/O group 0 contains nodes A and B
- I/O group 1 contains nodes C and D



Figure 4. A four-node SAN Volume Controller cluster with the redundant ac power switch feature

- 1** I/O group 0
- 2** SAN Volume Controller node A
- 3** 2145-1U uninterruptible power supply A
- 4** SAN Volume Controller node B
- 5** 2145-1U uninterruptible power supply B
- 6** I/O group 1
- 7** SAN Volume Controller node C
- 8** 2145-1U uninterruptible power supply C
- 9** SAN Volume Controller node D
- 10** 2145-1U uninterruptible power supply D
- 11** Redundant ac power switch 1
- 12** Redundant ac power switch 2

| **13** Site PDU X (C13 outlets)

| **14** Site PDU Y (C13 outlets)

| The site PDUs X and Y (**13** and **14**) are powered from two independent power
| sources.

| In this example, only two redundant ac power switch units are used, and each
| power switch powers one node in each I/O group. However, for maximum
| redundancy, use one redundant ac power switch to power each node in the cluster.

Chapter 3. SAN fabric and zoning overview

The *SAN fabric* is an area of the network that contains routers, gateways, hubs, and switches. In fibre-channel environments, *zoning* is the grouping of multiple ports to form a virtual, private, storage network.

SAN fabric overview

The SAN fabric is an area of the network that contains routers, gateways, hubs, and switches. A SAN is configured into a number of zones. A device using the SAN can only communicate with devices that are included in the same zones that it is in. A SAN Volume Controller cluster requires two distinct types of zones: a host zone and a disk zone.

In the host zone, the host systems can identify and address the SAN Volume Controller nodes. You can have more than one host zone. Generally, you create one host zone for each host type. In the disk zone, the SAN Volume Controller nodes identify the disk drives. Host systems cannot operate on the disk drives directly; all data transfer occurs through the SAN Volume Controller nodes. Figure 5 shows several host systems that are connected in a SAN fabric.




Figure 5. Example of a SAN Volume Controller cluster in a fabric

A cluster of SAN Volume Controller nodes is connected to the same fabric and presents virtual disks (VDisks) to the host systems. You create these VDisks from units of space within a managed disk (MDisk) group. An MDisk group is a collection of MDisks that are presented by the storage subsystems (RAID

controllers). The MDisk group provides a storage pool. You specify how each group is created, and you can combine MDisks from different manufacturers' controllers in the same MDisk group.

Note: Some operating systems cannot tolerate other operating systems in the same host zone, although you might have more than one host type in the SAN fabric. For example, you can have a SAN that contains one host that runs on an AIX® operating system and another host that runs on a Windows® operating system.

SAN fabric configuration

Ensure that you understand the rules and requirements when you are configuring the SAN fabric.

Table 1 provides terms and definitions that can guide your understanding of the rules and requirements.

Table 1. Configuration terms and definitions

Term	Definition
ISL hop	A hop on an interswitch 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. Note: The SAN Volume Controller puts its back-end traffic onto the same symmetrical network. The back-end traffic can vary 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.
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 and Global Mirror, significant distances might exist between the components of the local cluster and those of the remote cluster.

Table 1. Configuration terms and definitions (continued)

Term	Definition
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 and Global 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 might be significant distances between the components in the local cluster and those in the remote cluster. These components might be single-mode optical fibers 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.
Not valid configuration	The current SAN configuration is not correct. An attempted operation failed and generated an error code to indicate what caused it to become not valid. The most likely cause is that either a device has failed, or a device has been added to the SAN that has caused the configuration to be marked as not valid.
Unsupported configuration	A configuration that might operate successfully, but for which IBM does not guarantee the solution for problems that might occur. Usually this type of configuration does not create an error log entry.
Valid configuration	A configuration that consists of devices and connections that are identified as a valid and supported configuration. Neither of the following two conditions exist with the current configuration: <ul style="list-style-type: none"> • Not valid • Unsupported configuration
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 that connects other SAN fabric components. Generally these components might 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 that are configured to create a larger switched network. With this configuration, four or more switches are connected in a loop with some of the paths short circuiting the loop. An example of this configuration is four switches that are connected in a loop with ISLs for one of the diagonals.

SAN hardware configuration

Ensure that you are familiar with the configuration rules for fibre-channel switches. You must follow the configuration rules for fibre-channel switches to ensure that you have a valid configuration.

The SAN must contain only supported switches.

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

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

The SAN should consist of at least 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. Therefore, SANs with one fabric are considered to have a possible single point of failure.

Configurations with more than four SANs are not supported.

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 configuration that uses direct connections between the host and node or between the controller and node is not supported.

All backend storage must always and only be connected to SAN switches. Multiple connections are permitted from the redundant controllers of the backend storage to improve data bandwidth performance. It is not necessary to have a connection between each redundant disk controller system of the backend storage and each counterpart SAN. For example, in an IBM System Storage DS4000 configuration in which the IBM DS4000 contains two redundant controllers, only two controller minihubs are usually used. Controller A of the IBM DS4000 is connected to counterpart SAN A, and controller B of the IBM DS4000 is connected to counterpart SAN B. Any configuration that uses a direct connection between the host and the controller is not supported.

When you attach a node to a SAN fabric that contains core directors and edge switches, connect the node ports to the core directors and connect the host ports to the edge switches. In this type of 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. Any configuration that runs outside the manufacturer's configuration rules is not supported.

Mixing manufacturer switches in a single SAN fabric

Within an individual SAN fabric, switches must have the same manufacturer, with the following exceptions:

- IBM BladeCenter® products. The documentation that is provided with your BladeCenter unit has more information.
- Where one pair of counterpart fabrics (for example, Fabric A and Fabric B) provide a redundant SAN, different manufacturer's switches can be mixed in a SAN Volume Controller configuration, provided that each fabric contains only switches from a single manufacturer. Thus, the two counterpart SANs can have different manufacturer's switches.
- SAN Volume Controller supports interoperability between McData and Brocade products. For further details, refer to switch vendor documentation.

For additional information on BladeCenter support and other current interoperability information, see the following Web site:

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

Brocade core-edge fabrics

For Brocade core-edge configurations with greater than 64 hosts, you must follow these requirements:

SAN Volume Controller software level 4.1.1 or higher

Brocade core-edge fabrics that use the M14, M48, or B64 models can have up to 1024 hosts under the following conditions:

- M14, M48, B64 or other Brocade models can be used as edge switches; however, the SAN Volume Controller ports and backend storage must all be connected to the M14, M48, or B64 core switch.
- The M48 and B64 models must be running the firmware level 5.1.0c or higher.
- The M14 models must be running at the firmware level 5.0.5a or higher.

Fibre-channel switches and interswitch links

The SAN Volume Controller supports distance-extender technology, including DWDM (dense wavelength division multiplexing) and FCIP (Fibre Channel over IP) extenders, to increase the overall distance between local and remote clusters. If this extender technology involves a protocol conversion, the local and remote fabrics are regarded as independent fabrics, limited to three ISL hops each.

Note: Where multiple ISL hops are used between switches, follow the fabric manufacturer's recommendations for trunking.

With ISLs between nodes in the same cluster, the ISLs are considered a single point of failure. This is illustrated in Figure 6.




Figure 6. Fabric with ISL 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 fails for a period of time, and the node is not 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 7.

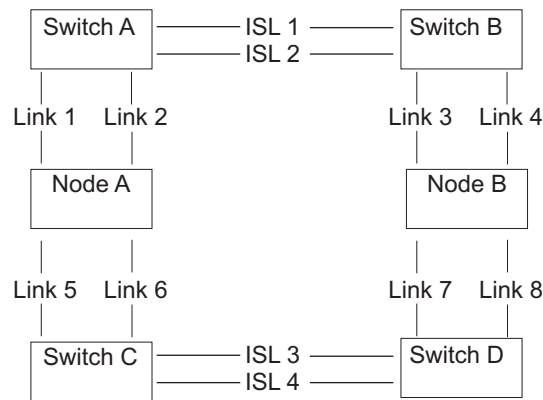


Figure 7. Fabric with ISL in a redundant configuration

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

ISL oversubscription

Perform a thorough SAN design analysis to avoid ISL congestion. Do not configure the SAN to use SAN Volume Controller to SAN Volume Controller traffic or SAN Volume Controller to storage subsystem traffic across ISLs. For host to SAN Volume Controller traffic, do not use an ISL oversubscription ratio that is greater than 7 to 1. Congestion on the ISLs can result in severe SAN Volume Controller performance degradation and I/O errors on the host.

When you calculate oversubscription, you must account for the speed of the links. For example, if the ISLs run at 4 Gbps and the host runs at 2 Gbps, calculate the port oversubscription as $7 \times (4/2)$. In this example, the oversubscription can be 14 ports for every ISL port.

Note: The SAN Volume Controller port speed is not used in the oversubscription calculation.

SAN Volume Controller in a SAN with director class switches

You can use director class switches within the SAN to connect large numbers of RAID controllers and hosts to a SAN Volume Controller cluster. Because director class switches provide internal redundancy, one director class switch can replace a SAN that uses multiple switches. However, the director class switch provides only network redundancy; it does not protect against physical damage (for example, flood or fire), which might destroy the entire function. A tiered network of smaller switches or a core-edge topology with multiple switches in the core can provide comprehensive redundancy and more protection against physical damage for a network in a wide area.

Zoning guidelines

Ensure that you are familiar with the zoning guidelines for controller and host zones.

Paths to hosts

- The number of paths through the network from the SAN Volume Controller nodes to a host must not exceed eight. Configurations in which this number is exceeded are not supported.
 - Each node has four ports and each I/O group has two nodes. Therefore, without any zoning, the number of paths to a VDisk would be $8 \times$ the number of host ports.
 - This rule exists to limit the number of paths that must be resolved by the multipathing device driver.

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.

Storage controller zones

Switch zones that contain controller ports must not have more than 40 ports. A configuration that exceeds 40 ports is not supported.

SAN Volume Controller zones

The switch fabric must be zoned so that the SAN Volume Controller nodes can see the back-end storage and the front-end host HBAs. Typically, the front-end host HBAs and the back-end storage are not in the same zone. The exception to this is where split host and split controller configuration is in use. All nodes in a 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 logs errors that request a repair action. This can occur if inappropriate zoning was applied to the fabric or if inappropriate LUN masking is used. This rule has important implications for back-end storage, such as an IBM DS4000, which imposes exclusive rules for mappings between HBA worldwide node names (WWNNs) and storage partitions.

It is possible to zone the switches in such a way that a SAN Volume Controller port is used solely for internode communication, or for communication to a host, or for communication to back-end storage. This is possible because each node contains four ports. Each node must still remain connected to the full SAN fabric. Zoning cannot be used to separate the SAN into two parts.

It is critical that you configure the controller and SAN so that a cluster cannot access LUs that a host or another cluster can also access. This can be arranged by controller LUN mapping and masking.

All nodes in a cluster must see at least one node port for all nodes that are in the cluster, but nodes cannot see node ports for nodes that belong to another cluster. You can have nodes that are not members of any cluster zoned to see all of the clusters. This allows you to add a node to the cluster in the event that you must replace a node.

With Metro Mirror and Global Mirror configurations, additional zones are required that contain only the local nodes and the remote nodes. It is valid for the local hosts to see the remote nodes or for the remote hosts to see the local nodes. Any zone that contains the local and the remote back-end storage and local nodes or remote nodes, or both, is not valid.

If a node can see another node through multiple paths, use zoning where possible to ensure that the node to node communication does not travel over an ISL. If a node can see a storage controller through multiple paths, use zoning to restrict communication to those paths that do not travel over ISLs.

Host zones

The configuration rules for host zones are different depending upon the number of hosts that will access the cluster. For smaller configurations of less than 64 hosts per cluster, the SAN Volume Controller supports a simple set of zoning rules which enable a small set of host zones to be created for different environments. For larger configurations of more than 64 hosts, the SAN Volume Controller supports a more restrictive set of host zoning rules.

Zoning that contains host HBAs must not contain either host HBAs in dissimilar hosts or dissimilar HBAs in the same host that are in separate zones. Dissimilar hosts means that the hosts are running different operating systems or are different hardware platforms; thus different levels of the same operating system are regarded as similar.

To obtain the best overall performance of the subsystem and to prevent overloading, the workload to each SAN Volume Controller port must be equal. This can typically involve zoning approximately the same number of host Fibre Channel ports to each SAN Volume Controller Fibre Channel port.

Clusters with less than 64 hosts

For clusters with less than 64 hosts attached, zones that contain host HBAs must contain no more than 40 initiators including the SAN Volume Controller ports that act as initiators. A configuration that exceeds 40 initiators is not supported. A valid zone can be 32 host ports plus 8 SAN Volume Controller ports. When it is possible, place each HBA port in a host that connects to a node into a separate zone. Include exactly one port from each node in the I/O groups that are associated with this host. This type of host zoning is not mandatory, but is preferred for smaller configurations.

Note: If the switch vendor recommends fewer ports per zone for a particular SAN, the rules that are imposed by the vendor takes precedence over the SAN Volume Controller rules.

To obtain the best performance from a host with multiple Fibre Channel ports, the zoning must ensure that each Fibre Channel port of a host is zoned with a different group of SAN Volume Controller ports.

Clusters with more than 64 hosts

Each HBA port must be in a separate zone and each zone must contain exactly one port from each SAN Volume Controller node in each I/O group that the host accesses.

Note: A host can be associated with more than one I/O group and therefore access VDisks from different I/O groups in a SAN. However, this reduces the maximum number of hosts that can be used in the SAN. For example, if the same host uses VDisks in two different I/O groups, this consumes one of the 256 hosts in each I/O group. If each host accesses VDisks in every I/O group, there can be only 256 hosts in the configuration.

Zoning examples

These examples describe ways for zoning a switch.

Example 1

Consider the SAN environment in the following example:

- Two nodes (nodes A and B)
- Nodes A and B each have four ports
 - Node A has ports A0, A1, A2, and A3
 - 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 2.

Table 2. 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

The following is an example configuration:

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.

Create the following host zones on switch X:

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.

Create the following host zones on switch Y:

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.

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.

Example 2

The following example describes a SAN environment that is similar to the previous example except for the addition of two hosts that have two ports each.

- Two nodes called A and B
- Nodes A and B have four ports each
 - Node A has ports A0, A1, A2, and A3
 - 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 the other two hosts have two ports each as described in Table 3.

Table 3. 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

The following is an example configuration:

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

Create the following host zones on switch X:

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.

Create the following host zones on switch Y:

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.

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.

Zoning considerations for Metro Mirror and Global Mirror

Ensure that you are familiar with the constraints for zoning a switch to support the Metro Mirror and Global Mirror features.

SAN configurations that use intracluster Metro Mirror and Global Mirror relationships do not require additional switch zones.

SAN configurations that use intercluster Metro Mirror and Global Mirror relationships require the following additional switch zoning considerations:

- The clusters must be zoned so that the nodes in each cluster can see the ports of the nodes in the other cluster.
- Use of interswitch link (ISL) trunking in a switched fabric.
- Use of redundant fabrics.

For intercluster Metro Mirror and Global Mirror relationships, you must perform the following steps to create the additional zones that are required:

1. Configure your SAN so that fibre-channel traffic can be passed between the two clusters. To configure the SAN this way, you can connect the clusters to the same SAN, merge the SANs, or use routing technologies.
2. Configure zoning to allow all nodes in the local fabric to communicate with all nodes in the remote fabric.

Note: If you are using McData Eclipse routers to route between two SANs, you can connect a maximum of 64 port pairs over a single iFCP link. For larger clusters, you must create several zones that contain at least one port from each node in both the local cluster and the remote cluster. For example, if you have two 8 node clusters and select one port from each node, there are 64 port pairs in each zone. For maximum connectivity, you must have four zones and 4 iFCP links. For smaller clusters, you can use less zones and iFCP links.

3. Optionally, modify the zoning so that the hosts that are visible to the local cluster can recognize the remote cluster. This allows a host to examine data in both the local and remote cluster.
4. Verify that cluster A cannot recognize any of the back-end storage that is owned by cluster B. Two clusters cannot share the same back-end storage devices.

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.

The following table provides a description of the ISL trunking and the extended fabric features:

Feature	Description
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 might 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 cannot be used together. 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. Therefore, do not use mixed mode operations.</p>
Extended fabric	<p>Extended fabric operation allocates extra buffer credits to a port. This is important over long links that are usually found in intercluster Metro Mirror operation. Because of the time that it takes for a frame to traverse the link, it is possible to have more frames in transmission at any instant in time than is 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"> • The Normal option is suitable for short links. • The Extended Normal option provides significantly better performance for the links up to 10 km long. <p>Note: The extended fabric license provides two extra options: Medium, 10 - 50 km and Long, 50 - 100 km. Do not use Medium and Long settings in the intercluster Metro Mirror links that are currently supported.</p>

Chapter 4. Configuration planning for the SAN Volume Controller

Ensure that you perform all the required and necessary planning tasks before you start to configure your SAN Volume Controller environment.

Planning the clusters

Determine the following information for clusters:

- The number of clusters and the number of node pairs (I/O groups). Each pair of nodes is the container for one or more virtual disks (VDisks)
- The number of hosts that you want to use
- The number of I/Os per second between the hosts and nodes

Planning the hosts

LUN masking allows the hosts to access specific logical units (LUs) within the disk controllers. Determine the following information for hosts:

- The worldwide port names (WWPNs) of the fibre-channel host bus adapter (HBA) ports in the hosts.
- The names to assign to the hosts.
- The VDIsks to assign to the hosts.

Planning the MDisks

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

Planning the managed disk groups

Determine the following information for MDisk groups:

- The types of backend controllers that you want to use.
- If you want to create VDIsks with the sequential policy, plan to create a separate MDisk group for these VDIsks or ensure that you create these VDIsks before creating VDIsks 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.
- Plan the extent size of the managed MDisk group. For example, a larger extent size increases the total amount of storage which the SAN Volume Controller can manage. A smaller extent size provides more fine-grained control of storage allocation.

Planning the VDIsks

An individual VDisk is a member of one managed disk group and one I/O group. The managed disk group defines which MDIsks provide the backend storage that makes up the VDisk. The I/O group defines which nodes provide I/O access to the VDisk. Before you create a VDisk, determine the following information:

- If the VDisk should be created in image mode from a managed disk that contains data that needs to be preserved.
- The name that you want to assign to the VDisk.
- The I/O group to which the VDisk will be assigned.
- The managed disk group to which the VDisk will be assigned. For example, different managed disk groups could have different performance characteristics depending on the storage that is contained by that managed disk group.
- The capacity of the VDisk.
- If you want to provide extra redundancy by mirroring the VDisk across managed disk groups. For example, you could use the VDisk Mirroring feature to provide redundancy across managed disk groups.
- If you want to create fully committed VDIsks or use space-efficient virtual disks.

SAN Volume Controller maximum configuration

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

See the following Web site for the latest maximum configuration support:

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

Configuration rules

Storage area network (SAN) configurations that contain SAN Volume Controller nodes can be configured in various ways.

A SAN configuration that contains SAN Volume Controller nodes must follow the rules for the following components:

- Storage subsystems
- Host bus adapters
- Nodes
- Fibre-channel switches
- Fabrics
- Zoning

Storage subsystem configuration rules

Follow these rules when you are planning the configuration of storage subsystems for use with SAN Volume Controller clusters.

See the following Web site for the latest support information:

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

All SAN Volume Controller nodes in 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 an IBM System Storage DS4000 series controller, which has exclusion rules that determine to which host bus adapter (HBA) worldwide node names (WWNNs) a storage partition can be mapped.

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.

You can configure certain storage controllers to safely share resources between the SAN Volume Controller cluster and direct attached hosts. This type of configuration is described as a split controller. In all cases, it is critical that you configure the controller and SAN so that the SAN Volume Controller cluster cannot access logical units (LUs) that a host or another SAN Volume Controller cluster can also access. This split controller configuration can be arranged by controller logical unit number (LUN) mapping and masking. If the split controller configuration is not guaranteed, data corruption can occur.

Besides a configuration where a controller is split between a SAN Volume Controller cluster and a host, the SAN Volume Controller cluster also supports configurations where a controller is split between two SAN Volume Controller clusters. In all cases, it is critical that you configure the controller and SAN so that the SAN Volume Controller cluster cannot access LUs that a host or another SAN Volume Controller cluster can also access. This can be arranged by controller LUN mapping and masking. If this is not guaranteed, data corruption can occur. Do not use this configuration because of the risk of data corruption.

Avoid configuring one storage subsystem device 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 cluster must be configured to manage only LUNs that are presented by supported disk controller systems. Operation with other devices is not supported.

Unsupported storage subsystem (generic device) rules

When a storage subsystem is detected on the SAN, the SAN Volume Controller attempts to recognize it using its Inquiry data. If the device is recognized as one of the explicitly supported storage models, the SAN Volume Controller uses error recovery programs that are potentially tailored to the known needs of the storage subsystem. If the device is not recognized, the SAN Volume Controller configures the device as a generic device. A generic device might not function correctly when it is addressed by a SAN Volume Controller cluster. In any event, the SAN Volume Controller cluster does not regard accessing a generic device as an error condition and, consequently, does not log an error. Managed disks (MDisks) that are presented by generic devices are not eligible to be used as quorum disks.

Split controller configurations rules

The SAN Volume Controller cluster is configured to manage LUs that are exported only by RAID controllers. Operation with other RAID controllers is illegal. While it is possible to use the SAN Volume Controller cluster to manage JBOD (just a bunch of disks) LUs that are presented by supported RAID controllers, the SAN Volume Controller cluster itself does not provide RAID functions, so these LUs are exposed to data loss in the event of a disk failure.

If a single RAID controller presents multiple LUs, either by having multiple RAID configured or by partitioning one or more RAID into multiple LUs, each LU can be owned by either the SAN Volume Controller cluster or a directly attached host.

Suitable LUN masking must be in place to ensure that LUs are not shared between SAN Volume Controller nodes and direct attached hosts.

In a split controller configuration, a RAID presents some of its LUs to a SAN Volume Controller cluster (which treats the LU as an MDisk) and the remaining LUs to another host. The SAN Volume Controller cluster presents virtual disks (VDisks) that are created from the MDisk to another host. There is no requirement for the multipathing driver for the two hosts to be the same. Figure 8 shows that the RAID controller is an IBM DS4000, with RDAC used for pathing on the directly attached host, and SDD used on the host that is attached with the SAN Volume Controller. Hosts can simultaneously access LUs that are provided by the SAN Volume Controller cluster and directly by the device.




Figure 8. Disk controller system shared between SAN Volume Controller node and a host

It is also possible to split a host so that it accesses some of its LUNs through the SAN Volume Controller cluster and some directly. In this case, the multipathing software that is used by the controller must be compatible with the SAN Volume Controller multipathing software. Figure 9 on page 49 is a supported configuration because the same multipathing driver is used for both directly accessed LUNs and VDIs.




Figure 9. IBM System Storage DS8000 LUs accessed directly with a SAN Volume Controller node

In the case where the RAID controller uses multipathing software that is compatible with SAN Volume Controller multipathing software (see Figure 10 on page 50), it is possible to configure a system where some LUNs are mapped directly to the host and others are accessed through the SAN Volume Controller. An IBM TotalStorage® Enterprise Storage Server® (ESS) that uses the same multipathing driver as a SAN Volume Controller node is one example.




Figure 10. IBM DS4000 direct connection with a SAN Volume Controller node on one host

Host bus adapter configuration rules

Ensure that you are familiar with the configuration rules for host bus adapters (HBAs). You must abide by the configuration rules for HBAs to ensure that you have a valid configuration.

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.ibm.com/storage/support/2145>

Operation with other HBAs is not supported.

The SAN Volume Controller does not specify the number of host fibre-channel ports or HBAs that a host or a partition of a host can have. The number of host fibre-channel ports or HBAs are specified by the host multipathing device driver. The SAN Volume Controller supports this number; however it is subject to the configuration rules for the SAN Volume Controller. To obtain optimal performance and to prevent overloading, the workload to each SAN Volume Controller port

must be equal. You can achieve an even workload by zoning approximately the same number of host fibre-channel ports to each SAN Volume Controller fibre-channel port.

You can attach the SAN Volume Controller to open-systems hosts that use the small computer system interface-fibre channel protocol (SCSI-FCP). You can also attach the SAN Volume Controller to iSCSI (small computer system interface over internet protocol) hosts using FCIP ports in your SAN fabric. iSCSI hosts are supported only in non-failover configurations.

Node configuration rules

You must follow the configuration rules for SAN Volume Controller nodes to ensure that you have a valid configuration.

Host bus adapters and nodes

SAN Volume Controller 2145-4F2 and SAN Volume Controller 2145-8F2 nodes contain two 2-port host bus adapters (HBAs). If one HBA fails, the node operates in degraded mode. If an HBA is physically removed, the configuration is not supported.

SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8G4 nodes contain one 4-port HBA.

I/O groups

Nodes must always be used in pairs called I/O groups. SAN Volume Controller 2145-4F2, SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, and SAN Volume Controller 2145-8G4 nodes can be in the same I/O group. If a node fails or is removed from the configuration, the remaining node in the I/O group operates in a degraded mode, but the configuration is still valid.

VDisks

Each node presents a virtual disk (VDisk) to the SAN through four ports. Each VDisk is accessible from the two nodes in an I/O group. Each host HBA port can recognize up to eight paths to each logical unit (LU) that is presented by the node. The hosts must run a multipathing device driver before the multiple paths can resolve to a single device.

Optical connections

Valid optical connections are based on the fabric rules that the manufacturers impose for the following connection methods:

- Host to a switch
- Backend to a switch
- Interswitch links (ISLs)

Short-wave optical fiber connections must be used between a node and its switches. Clusters that use the intercluster Metro Mirror or Global Mirror features can use short or long wave optical fiber connections between the switches, or they can use distance-extender technology that is supported by the switch manufacturer.

The number of paths through the network from the node to a host must not exceed eight. Configurations in which this number is exceeded are not supported.

Each node has four ports and each I/O group has two nodes. Therefore, without any zoning, the number of paths to a VDisk is eight × the number of host ports.

Ethernet connection

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

Physical location

The physical distance between SAN Volume Controller nodes in the same cluster is limited to 100 meters due to connectivity requirements and servicing requirements. Several of the SAN Volume Controller service actions in problem situations require that the manipulations be done to both SAN Volume Controller nodes within an I/O group or a cluster within one minute of each other. Set up your cluster environment to enable IBM service personnel to easily perform actions that are almost simultaneous in the required timeframe.

A SAN Volume Controller node must be in the same rack as the uninterruptible power supply from which it is supplied.

Fibre-channel connection

The SAN Volume Controller supports short-wave, small-form factor pluggable (SFP) transceivers (850 nm with 50 µm or 62.5 µm multimode cables) between the SAN Volume Controller nodes and the switch to which they are connected. The transceivers can run at up to 500 m and are limited by the pulse spreading that is caused by the multimode nature of the transmission.

To avoid communication between nodes that are being routed across interswitch links (ISLs), connect all SAN Volume Controller nodes to the same fibre-channel switches.

No ISL hops are permitted among the SAN Volume Controller nodes within the same I/O group. However, one ISL hop is permitted among SAN Volume Controller nodes that are in the same cluster though different I/O groups. If your configuration requires more than one ISL hop for SAN Volume Controller nodes that are in the same cluster but in different I/O groups, contact your IBM service representative.

To avoid communication between nodes and storage subsystems that are being routed across ISLs, connect all storage subsystems to the same fibre-channel switches as the SAN Volume Controller nodes. One ISL hop between the SAN Volume Controller nodes and the storage controllers is permitted. If your configuration requires more than one ISL, contact your IBM service representative.

In larger configurations, it is common to have ISLs between host systems and the SAN Volume Controller nodes.

Port speed

You can change the operational port speed for SAN Volume Controller 2145-4F2 and SAN Volume Controller 2145-8F2 nodes to 1 Gbps or 2 Gbps. However, the optical fiber connections between the fibre-channel switches and all SAN Volume Controller 2145-4F2 and SAN Volume Controller 2145-8F2 nodes in a cluster must run at the same speed. The fibre-channel ports on SAN Volume Controller

2145-8F4 and SAN Volume Controller 2145-8G4 nodes auto negotiate the operational port speed independently, which allows these nodes to operate at different speeds. SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8G4 nodes can operate at 1 Gbps, 2 Gbps or 4 Gbps. If these nodes are connected to a 4 Gbps capable switch, the port attempts to operate at 4 Gbps; however, if there is a large number of link error rates, the adapter negotiates a lower speed.

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 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: WebKing v5.5 and Window-Eyes v5.5.
- You can operate all features using the keyboard instead of the mouse.
- You can change the initial delay and repeat rate of the up and down buttons to two seconds when you use the front panel of the SAN Volume Controller to set or change an IPv4 address. This feature is documented in the applicable sections of the SAN Volume Controller publications.

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 at the following Web site:

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

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.*

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

*IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS 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.

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

*IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.*

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

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

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

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

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

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

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

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

Trademarks

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

- AIX
- BladeCenter
- Enterprise Storage Server
- FlashCopy
- IBM
- IBM eServer
- IBM TotalStorage
- IBM System Storage
- System p5
- System z9
- System Storage
- 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.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

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 to indicate special notices.

The notices throughout the SAN Volume Controller documentation and in the *IBM Systems Safety Notices* document follow specific guidelines for their content.

The following notices are used throughout this library to convey 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 the situation in which damage might 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.

The caution and danger notices that appear in this document are also in the multilingual *IBM Systems Safety Notices* document. Use the reference numbers in parentheses; for example (1), at the end of each notice to find the matching translated notice.

Glossary

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

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

The following cross-references are used in this glossary:

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

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

See also

Refers the reader to one or more related terms.

Contrast with

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

Numerics

2145 A hardware machine type for the IBM System Storage SAN Volume Controller. Models of the SAN Volume Controller are expressed as the number 2145 followed by "-xxx", such as 2145-8G4. Hardware models for the 2145 include 2145-4F2, 2145-8F2, 2145-8F4, and 2145-8G4.

A

access mode

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

Address Resolution Protocol (ARP)

A protocol that dynamically maps an IP address to a network adapter address in a local area network.

agent code

An open-systems standard that interprets Common Information Model (CIM) requests and responses as they transfer between the client application and the device.

application server

A host that is attached to the storage area network (SAN) and that runs applications.

ARP See *Address Resolution Protocol*.

array An ordered collection, or group, of physical storage devices that are used to define logical volumes or devices.

association

A class that contains two references that define a relationship between two referenced objects.

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*.

availability

The ability of a system to continue working, with perhaps a decrease in performance, after individual components fail.

B**bandwidth**

The range of frequencies an electronic system can transmit or receive. The greater the bandwidth of a system, the more information the system can transfer in a given period of time.

bitmap

A coded representation in which each bit, or group of bits, represents or corresponds to an item; for example, a configuration of bits in main storage in which each bit indicates whether a peripheral device or a storage block is available or in which each group of bits corresponds to one pixel of a display image.

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*.

Boolean

Pertaining to the processes used in the algebra formulated by George Boole.

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.

Call Home

In SAN Volume Controller, a communication service that sends data and

event notifications to a service provider. The machine can use this link to place a call to IBM or to another service provider when service is required.

cascading

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

CIM See *Common Information Model*.

CIM object manager (CIMOM)

The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application. It then directs the requests to the appropriate component or service provider.

CIMOM

See *CIM object manager*.

class The definition of an object within a specific hierarchy. A class can have properties and methods and can serve as the target of an association.

CLI See *command line interface*.

client A computer system or process that requests a service of another computer system or process that is typically referred to as a server. Multiple clients can share access to a common server.

client application

A storage management program that initiates Common Information Model (CIM) requests to the CIM agent for the device.

cluster

In SAN Volume Controller, up to four pairs of nodes that provide a single configuration and service 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.

concurrent maintenance

Service that is performed on a unit while it is operational.

In SAN Volume Controller, the ability for one node in the cluster to be turned off for maintenance without interrupting access to the VDisk data provided by the cluster.

configuration node

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

connected

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

consistency group

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

consistent copy

In a Metro or 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.

consistent-stopped

In a Global Mirror relationship, the state that occurs when the secondary virtual disk (VDisk) contains a consistent image, but the image might be out-of-date with respect to the primary VDisk. This state can happen when a relationship was in the consistent-synchronized state when an error occurred that forced a freeze of the consistency group. This state can also happen when a relationship is created with the create-consistent flag set to TRUE.

consistent-synchronized

In a Global Mirror relationship, the status condition that occurs when the primary virtual disk (VDisk) is accessible for read and write I/O operations. The secondary VDisk is accessible for read-only I/O operations. See also *primary virtual disk* and *secondary virtual disk*.

container

A data storage location; for example, a file, directory, or device.

A software object that holds or organizes other software objects or entities.

contingency capacity

Initially, a fixed amount of unused real capacity that is maintained on a space-efficient virtual disk that is configured to automatically expand its real capacity. It is also the difference between the used capacity and the new real capacity when the real capacity is changed manually.

copied

In a FlashCopy mapping, 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.

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.

Copy Services

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

counterpart SAN

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

cross-volume consistency

In SAN Volume Controller, a consistency group property that guarantees consistency between virtual disks when an application issues dependent write operations that span multiple virtual disks.

D

data migration

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

degraded

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

dense wavelength division multiplexing (DWDM)

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

dependent write operations

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

destage

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

device In the CIM Agent, the storage server that processes and hosts client application requests.

IBM definition: A piece of equipment that is used with the computer and does not generally interact directly with the system, but is controlled by a controller.

HP definition: In its physical form, a magnetic disk that can be attached to a SCSI bus. The term is also used to indicate a physical device that has been made part of a controller configuration; that is, a physical device that is known to the controller. Units (virtual disks) can be created from devices after the devices have been made known to the controller.

device provider

A device-specific handler that serves as a plug-in for the Common Information Model (CIM); that is, the CIM object manager (CIMOM) uses the handler to interface with the device.

directed maintenance procedures

The set of maintenance procedures that can be run for a cluster. These procedures are run from within the SAN Volume Controller application and are documented in the *IBM System Storage SAN Volume Controller: Service Guide*.

disconnected

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

discovery

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

disk controller

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

disk drive

A disk-based, nonvolatile, storage medium.

disk zone

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

Distributed Management Task Force (DMTF)

An organization that defines standards for the management of distributed systems. See also *Common Information Model*.

DMP See *directed maintenance procedures*.

DMTF

See *Distributed Management Task Force*.

domain name server

In the Internet suite of protocols, a server program that supplies name-to-address conversion by mapping domain names to IP addresses.

DRAM

See *dynamic random access memory*.

DWDM

See *dense wavelength division multiplexing*.

dynamic random access memory (DRAM)

A storage in which the cells require repetitive application of control signals to retain stored data.

E

EC See *engineering change*.

EIA See *Electronic Industries Alliance*.

Electronic Industries Alliance (EIA)

An alliance of four trade associations: The Electronic Components, Assemblies & Materials Association (ECA); the Government Electronics and Information Technology Association (GEIA); the JEDEC Solid State Technology Association (JEDEC); and the Telecommunications Industry Association (TIA). Prior to 1998, EIA was the Electronic Industries Association and the group dates back to 1924.

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

engineering change (EC)

A correction for a defect of hardware or software that is applied to a product.

error code

A value that identifies an error condition.

ESS See *IBM TotalStorage Enterprise Storage Server*.

exclude

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

excluded

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

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

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*.

fabric port (F_port)

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

failover

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

FCIP See *Fibre Channel over IP*.

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.

fibre-channel extender

A device that extends a fibre-channel link over a greater distance than is supported by the standard, usually a number of miles or kilometers. Devices must be deployed in pairs at each end of a link.

Fibre Channel over IP (FCIP)

A network storage technology that combines the features of the Fibre Channel Protocol and the Internet Protocol (IP) to connect distributed SANs over large distances.

Fibre Channel Protocol (FCP)

A protocol that is used in fibre-channel communications with five layers that define how fibre-channel ports interact through their physical links to communicate with other ports.

field replaceable unit (FRU)

An assembly that is replaced in its entirety when any one of its components fails. An IBM service representative performs the replacement. In some cases, a field replaceable unit might contain other field replaceable units.

FlashCopy mapping

A relationship between two virtual disks.

FlashCopy relationship

See *FlashCopy mapping*.

FlashCopy service

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

F_port See *fabric port*.

FRU See *field replaceable unit*.

G

gateway

An entity that operates above the link layer and translates, when required, the interface and protocol used by one network into those used by another distinct network.

GB See *gigabyte*.

GBIC See *gigabit interface converter*.

gigabit interface converter (GBIC)

An interface module that converts the light stream from a fibre-channel cable into electronic signals for use by the network interface card.

gigabyte (GB)

In decimal notation, 1 073 741 824 bytes.

Global Mirror

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

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

graphical user interface (GUI)

A type of computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution graphics, pointing devices, menu bars and other menus, overlapping windows, icons and the object-action relationship.

GUI See *graphical user interface*.

H

hardcoded

Pertaining to software instructions that are statically encoded and not intended to be altered.

HBA See *host bus adapter*.

HLUN

See *virtual disk*.

hop One segment of a transmission path between adjacent nodes in a routed network.

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

host bus adapter (HBA)

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

host ID

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

host zone

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

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

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 System Storage Productivity Center (SSPC)

An integrated hardware and software solution that provides a single point of entry for managing SAN Volume Controller clusters, IBM System Storage DS8000 systems, and other components of a data storage infrastructure.

IBM TotalStorage Enterprise Storage Server (ESS)

An IBM product that provides an intelligent disk-storage subsystem across an enterprise.

ID See *identifier*.

identifier (ID)

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

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

idling The status of a pair of virtual disks (VDisks) that have a defined copy relationship for which no copy activity has yet been started.

In a Metro or Global Mirror relationship, the state that indicates that the master virtual disks (VDisks) and auxiliary VDIsks are operating in the primary role. Consequently, both VDIsks are accessible for write I/O operations.

idling-disconnected

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

illegal configuration

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

image mode

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

image VDisk

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

IML See *initial microcode load*.

inconsistent

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

inconsistent-copying

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

inconsistent-disconnected

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

inconsistent-stopped

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

indication

An object representation of an event.

initial microcode load (IML)

In SAN Volume Controller, the process by which the run-time code and data for a node are loaded into memory and initialized.

initiator

The system component that originates an I/O command over an I/O bus or network. I/O adapters, network interface cards, and intelligent controller device I/O bus control ASICs are typical initiators. (S) See also *logical unit number*.

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.

instance

An individual object that is a member of some class. In object-oriented programming, an object is created by instantiating a class.

integrity

The ability of a system to either return only correct data or respond that it cannot return correct data.

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. IPv4 is the dominant network layer protocol on the Internet, and IPv6 is designated as its successor. IPv6 provides a much larger address space, which enables greater flexibility in assigning addresses and simplifies routing and renumbering.

interswitch link (ISL)

The physical connection that carries a protocol for interconnecting multiple routers and switches in a storage area network.

I/O See *input/output*.

I/O group

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

I/O throttling rate

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

IP See *Internet Protocol*.

IP address

The unique 32-bit address that specifies the location of each device or workstation in the Internet. For example, 9.67.97.103 is an IP address.

ISL See *interswitch link*.

ISL hop

A hop on an interswitch link (ISL). Considering all pairs of node ports (N-ports) in a fabric and measuring distance only in terms of interswitch 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.

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

LBA See *logical block address*.

least recently used (LRU)

An algorithm used to identify and make available the cache space that contains the least-recently used data.

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.

local/remote fabric interconnect

The storage area network (SAN) components that are used to connect the local and remote fabrics together.

logical block address (LBA)

The block number on a disk.

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)

longitudinal redundancy check (LRC)

A method of error checking during data transfer that involves checking parity.

LRC See *longitudinal redundancy check*.

LRU See *least recently used*.

LU See *logical unit*.

LUN See *logical unit number*.

LUN masking

A process that allows or prevents I/O to the disk drives through the host-bus-adaptor (HBA) device or operating-system device driver.

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 console

A single point from which to manage the IBM System Storage SAN Volume Controller. For SAN Volume Controller version 4.2.1 and earlier, the master console was purchased either as software that was installed and configured on a server or as a hardware platform with preinstalled operating system and master console software. See *IBM System Storage Productivity Center*.

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*.

MB See *megabyte*.

MDisk

See *managed disk*.

megabyte (MB)

In decimal notation, 1 048 576 bytes.

mesh configuration

A network that contains a number of small SAN switches configured to create a larger switched network. With this configuration, four or more switches are connected together in a loop with some of the paths short circuiting the loop. An example of this configuration is to have four switches connected together in a loop with ISLs for one of the diagonals.

method

A way to implement a function on a class.

Metro Mirror

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

MIB See *Management Information Base*.

migration

See *data migration*.

mirrored virtual disk

A virtual disk (VDisk) with two VDisk copies.

mirrorset

IBM definition: See *RAID-1*.

HP definition: A RAID storageset of two or more physical disks that maintain a complete and independent copy of the data from the virtual disk. This type of storageset has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storagesets are referred to as mirrorsets.

N**namespace**

The scope within which a Common Information Model (CIM) schema applies.

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

node name

A name identifier associated with a node. (SNIA)

node port (N_port)

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

node rescue

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

non-RAID

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

N_port

See *node port*.

O

object In object-oriented design or programming, a concrete realization of a class that consists of data and the operations associated with that data.

object model

A representation, such as a diagram, of objects in a given system. Using symbols similar to standard flowchart symbols, an object model depicts the classes the objects belong to, their associations with each other, the attributes that make them unique, and the operations that the objects can perform and that can be performed on them.

object name

An object that consists of a namespace path and a model path. The namespace path provides access to the Common Information Model (CIM) implementation managed by the CIM Agent, and the model path provides navigation within the implementation.

object path

An object that consists of a namespace path and a model path. The namespace path provides access to the Common Information Model (CIM) implementation managed by the CIM Agent, and the model path provides navigation within the implementation.

offline

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

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

operating set

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

overallocated volume

See *space-efficient virtual disk*.

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 interswitch links (ISLs), where more than one ISL is connected in parallel between these switches. This definition assumes a symmetrical network and a specific workload that is applied equally from all initiators and sent equally to all targets. See also *symmetrical network*.

P

partition

IBM definition: A logical division of storage on a fixed disk.

HP definition: A logical division of a container represented to the host as a logical unit.

partner node

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

partnership

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

paused

In SAN Volume Controller, the process by which the cache component quiesces all ongoing I/O activity below the cache layer.

pend To cause to wait for an event.

petabyte (PB)

In decimal notation, 1 125 899 906 842 624 bytes.

PDU See *power distribution unit*.

PLUN See *managed disk*.

point-in-time copy

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

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

port ID

An identifier associated with a port.

power distribution unit (PDU)

A device that distributes electrical power to multiple devices in the rack. It typically is rack-mounted and provides circuit breakers and transient voltage suppression.

power-on self-test

A diagnostic test that servers or computers run when they are turned on.

prepared

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

preparing

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

primary virtual disk

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

property

In the Common Information Model (CIM), an attribute that is used to characterize instances of a class.

PuTTY

A client program that allows you to run remote sessions on your computer through specific network protocols, such as SSH, Telnet, and Rlogin.

Q**qualifier**

A value that provides additional information about a class, association, indication, method, method parameter, instance, property, or reference.

quorum

A set of nodes that operates as a cluster. Each node has a connection to every other node in the cluster. If a connection failure causes the cluster to split into two or more groups of nodes that have full connection within the group, the quorum is the group that is selected to operate as the cluster.

Typically, this is the larger group of nodes, but the quorum disk serves as a tiebreaker if the groups are the same size.

queue depth

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

quorum disk

A managed disk (MDisk) that contains a reserved area that is used exclusively for cluster management. The quorum disk is accessed in the event that it is necessary to determine which half of the cluster continues to read and write data.

quorum index

A number that can be either: 0, 1 or 2

R

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

RAID See *redundant array of independent disks*.

RAID 0

IBM definition: RAID 0 allows a number of disk drives to be combined and presented as one large disk. RAID 0 does not provide any data redundancy. If one drive fails, all data is lost.

HP definition: A RAID storage set that stripes data across an array of disk drives. A single logical disk spans multiple physical disks, allowing parallel data processing for increased I/O performance. While the performance characteristics of RAID level 0 is excellent, this RAID level is the only one that does not provide redundancy. Raid level 0 storage sets are referred to as stripe sets.

RAID 1

SNIA dictionary definition: A form of storage array in which two or more identical copies of data are maintained on separate media. (S)

IBM definition: A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirror set.

HP definition: See *mirror set*.

RAID 5

SNIA definition: A form of parity RAID in which the disks operate independently, the data strip size is no smaller than the exported block size, and parity check data is distributed across the array's disks. (S)

IBM definition: See the SNIA definition.

HP definition: A specially developed RAID storage set that stripes data and parity across three or more members in a disk array. A RAID set combines the best characteristics of RAID level 3 and RAID level 5. A RAID set is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAID set is sometimes called parity RAID. RAID level 3/5 storage sets are referred to as RAID sets.

RAID 10

A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set.

real capacity

The amount of storage that is allocated to a virtual disk copy from a managed disk group.

redundant ac power switch

A device that provides input power redundancy by attaching a SAN Volume Controller to two independent power sources. If the main source becomes unavailable, the redundant ac power switch automatically provides power from a secondary (backup) source. When power is restored, the redundant ac power switch automatically changes back to the main power source.

redundant array of independent disks (RAID)

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*.

reference

A pointer to another instance that defines the role and scope of an object in an association.

rejected

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

relationship

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

reliability

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

remote fabric

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

roles

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

S

SAN See *storage area network*.

SAN Volume Controller fibre-channel port fan in

The number of hosts that can see any one SAN Volume Controller port.

SATA See *Serial Advanced Technology Attachment*.

schema

A group of object classes defined for and applicable to a single namespace. Within the CIM Agent, the supported schemas are the ones that are loaded through the managed object format (MOF).

SCSI See *Small Computer Systems Interface*.

SCSI back-end layer

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

SCSI front-end layer

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

SDD See *subsystem device driver (SDD)*.

secondary virtual disk

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

Secure Shell (SSH)

A program to log in to another computer over a network, to run commands in a remote machine, and to move files from one machine to another.

Secure Sockets Layer (SSL)

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

sequential VDisk

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

Serial Advanced Technology Attachment (SATA)

The evolution of the ATA interface from a parallel bus to serial connection architecture. (S)

Serial ATA

See *Serial Advanced Technology Attachment*.

server In a network, the hardware or software that provides facilities to other stations; for example, a file server, a printer server, a mail server. The station making the request of the server is usually called the client.

Service Location Protocol (SLP)

In the Internet suite of protocols, a protocol that identifies and uses network hosts without having to designate a specific network host name.

fibres-channel SFP connector

See *small form-factor pluggable connector*.

Simple Mail Transfer Protocol (SMTP)

An Internet application protocol for transferring mail among users of the

Internet. SMTP specifies the mail exchange sequences and message format. It assumes that the Transmission Control Protocol (TCP) is the underlying protocol.

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).

SLP See *Service Location Protocol*.

Small Computer System Interface (SCSI)

A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

small form-factor pluggable (SFP) connector

A compact optical transceiver that provides the optical interface to a fibre-channel cable.

SMI-S See *Storage Management Initiative Specification*.

SMTP See *Simple Mail Transfer Protocol*.

SNIA See *Storage Networking Industry Association*.

SNMP

See *Simple Network Management Protocol*.

space-efficient VDisk

See *space-efficient virtual disk*.

space-efficient virtual disk

A virtual disk that has different virtual capacities and real capacities.

SSH See *Secure Shell*.

SSPC See *IBM System Storage Productivity Center (SSPC)*.

SSL See *Secure Sockets Layer*.

stand-alone relationship

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

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

stopped

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

storage area network (SAN)

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

Storage Management Initiative Specification (SMI-S)

A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface that allows storage management systems to identify, classify, monitor, and control

physical and logical resources in a storage area network. The interface is intended as a solution that integrates the various devices to be managed in a storage area network (SAN) and the tools used to manage them.

Storage Networking Industry Association (SNIA)

An association of producers and consumers of storage networking products whose goal is to further storage networking technology and applications. See www.snia.org.

striped

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

stripeset

See *RAID 0*.

subsystem device driver (SDD)

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

superuser authority

Can issue any command-line interface (CLI) command. A superuser can view and work with the following panels: View users, Add cluster, Remove cluster, Add users, and Modify users. Only one Superuser role is available.

suspended

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

switch

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

symmetrical network

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

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

system

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

T

terabyte

In decimal notation, 1 099 511 628 000 bytes.

thinly provisioned volume

See *space-efficient virtual disk*.

topology

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

trigger

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

U

UID See *unique identifier*.

unconfigured mode

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

uninterruptible power supply

A device that is 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.

unique identifier (UID)

An identifier that is assigned to storage system logical units when they are created. It is used to identify the logical unit regardless of the logical unit number (LUN), status of the logical unit, or whether alternate paths exist to the same device. Typically, a UID is only used once.

unmanaged

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

V

valid configuration

A configuration that is supported.

VDisk See *virtual disk (VDisk)*.

VDisk copy

See *virtual disk copy*.

virtual capacity

The amount of storage that is available to a server on a virtual disk (VDisk) copy. In a space-efficient virtual disk, the virtual capacity can be different from the real capacity. In a standard virtual disk, the virtual capacity and real capacity are the same.

virtual disk copy

A physical copy of the data that is stored on a virtual disk (VDisk). Mirrored VDIs have two such copies. Nonmirrored VDIs have one copy.

virtual disk (VDisk)

A device that host systems in a storage area network (SAN) recognize as a Small Computer System Interface (SCSI) disk.

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.

virtual storage area network (VSAN)

A fabric within the SAN.

vital product data (VPD)

Information that uniquely defines system, hardware, software, and microcode elements of a processing system.

VLUN See *managed disk*.

VPD See *vital product data*.

VSAN See *virtual storage area network*.

W**WBEM**

See *Web-Based Enterprise Management*.

Web-Based Enterprise Management (WBEM)

A tiered, enterprise-management architecture that was developed by the Distributed Management Task Force (DMTF). This architecture provides the management design framework that consists of devices, device providers, the object manager, and the messaging protocol for the communication between client applications and the object manager.

worldwide node name (WWNN)

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

worldwide port name (WWPN)

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

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

Z**zoning**

In fibre-channel environments, the grouping of multiple ports to form a

virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones.

Index

Numerics

- 2145-1U uninterruptible power supply
 - environment 22
 - power cables 22
- 2145-8G4 node
 - features 5
 - SAN Volume Controller 19

A

- ac power switch, cabling 30
- accessibility
 - keyboard 55
 - repeat rate of up and down buttons 55
 - shortcut keys 55
- application programming interface 6
- Assist On-site remote service 8

B

- Brocade
 - core-edge fabrics 37
 - switch ports 36
- browsers
 - / see also Web browsers 25

C

- cables
 - 2145-1U uninterruptible power supply power requirements 22
- Call Home 9
- capacity
 - real 14
 - virtual 14
- charts and tables
 - cable connection table 27
 - configuration data table 29
 - for planning 26
 - hardware location chart 26
 - redundant ac power switch connection chart 29
- CIM (Common Information Model) 6
- circuit breakers
 - requirements
 - SAN Volume Controller 2145-8G4 20
- clusters
 - Call Home e-mail 9
 - configuration backup overview 16
 - high availability 3
 - operation 11
 - overview 10
 - state 10
- Common Information Model (CIM) 6
 - configuration
 - maximum sizes 46
 - mesh 34

- configuration (*continued*)
 - planning 29
 - rules 34
- configuring
 - nodes 51
 - SAN 46
 - SAN Volume Controller 51
 - switches 35
- connections 25
- console
 - SAN Volume Controller
 - master console 5
 - user interface 5
- controller
 - configuration rules 46
- controllers
 - disk 13
 - zoning 39
- Copy Services
 - FlashCopy 16
 - Global Mirror 17
 - Metro Mirror 17
 - overview 16
 - zoning for Metro Mirror and Global Mirror 43
- copying
 - virtual disks 14

D

- data migration
 - image mode virtual disks 16
- disk controllers
 - overview 13

E

- e-mail
 - Call Home 9
 - inventory information 10
- event notification 9
- examples
 - redundant ac power switch cabling 30
 - SAN environments 41
 - SAN Volume Controller cluster in a SAN fabric 33

F

- fabric, SAN 33
- fibre-channel switches 35
- FlashCopy
 - definition 67
 - overview 16

G

- Global Mirror
 - overview 17
 - zoning considerations 43
- guidelines
 - zoning 39

H

- hardware
 - location chart 26
- HBAs (host bus adapters)
 - configuration 50
 - node 51
- high availability
 - cluster 3
- host bus adapters (HBAs)
 - configuration 50
 - node 51
- hosts
 - supported 5
 - zoning 39

I

- I/O groups 12, 51
- IBM System Storage Productivity Center 7
 - environment 25
- information
 - center xi
- installation
 - charts and tables 26
 - planning for SAN Volume Controller 19
- interfaces 5
- interswitch link (ISL)
 - congestion 38
 - maximum hop count 37
 - oversubscription 38
- inventory information 9

K

- keyboard 55

L

- legal notices 57

M

- maximum configuration 46
- mesh configuration 34
- Metro Mirror
 - overview 17
 - zoning considerations 43
- migration 15

- mirroring
 - virtual disks 14
- multipathing software 5

N

- nodes
 - configuration 51
 - host bus adapters (HBAs) 51
 - virtual disks (VDisks) 51
- notifications
 - inventory information 10
 - sending 9

O

- object descriptions in SAN Volume Controller environment 2
- operating over long distances 43
- oversubscription 34
- overview
 - Copy Services features 16
 - disk controllers 13
 - redundant ac power switch 6
 - SAN fabric 33
 - SAN Volume Controller 1
 - zoning 41

P

- PDU (power distribution unit) 22
- physical characteristics
 - 2145-1U uninterruptible power supply 22
 - redundant ac power switch 24
- planning
 - IBM System Storage Productivity Center 25
 - planning for SAN Volume Controller
 - configuration 45
 - configuration data table 29
 - installation 19
 - physical configuration 26
 - port speed 52
 - ports 25
 - power
 - cables
 - 2145-1U uninterruptible power supply 22
 - country or region 22
 - distribution unit (PDU) 22
 - requirements
 - SAN Volume Controller 2145-8G4 20
 - preparing
 - uninterruptible power supply environment 21
 - PuTTY 9

Q

- quorum disks
 - cluster operation 11

R

- redundant ac power switch
 - cabling 30
 - connection chart 29
 - environment preparation 24
 - examples 30
 - overview 6
 - power cables 22
 - specifications 24
- related information xi
- remote service 8
- requirements
 - 2145-1U uninterruptible power supply 6
 - ac voltage 19, 20
 - circuit breakers 20
 - electrical 19, 20
 - power 20
 - power cables 22
 - redundant ac power switch 22
 - SAN Volume Controller 2145-8G4 19
 - Web browsers 25

S

- SAN (storage area network)
 - configuring 46
 - fabric overview 33
 - I/O groups 12
- SAN fabric
 - configuring 34
- SAN Volume Controller
 - configuring nodes 51
 - Console
 - user interfaces 5
 - Web browser requirements 25
 - features 5
 - hardware 1
 - minimum requirements 5
 - object descriptions 2
 - overview 1
 - ports and connections 25
 - software
 - overview 1
 - virtualization 2
- SAN Volume Controller 2145-8G4
 - additional space requirements 21
 - air temperature without redundant ac power 20
 - circuit breaker requirements 20
 - dimensions and weight 21
 - heat output of node 21
 - humidity with redundant ac power 20
 - humidity without redundant ac power 20
 - input-voltage requirements 19
 - nodes
 - heat output 21
 - power requirements for each node 20
 - product characteristics 19
 - requirements 19
 - specifications 19
 - temperature with redundant ac power 20

- SAN Volume Controller 2145-8G4
 - (continued)
 - weight and dimensions 21
- SDD (subsystem device driver) 5
- Secure Shell (SSH)
 - overview 9
 - PuTTY 9
- service
 - remote through Assist On-site 8
- shortcut keys 55
- site requirements
 - connections 25
 - ports 25
- SNMP traps 9
- software
 - multipathing 5
 - overview 1
- space requirements
 - SAN Volume Controller 2145-8G4 21
- specifications
 - redundant ac power switch 24
- SSH (Secure Shell)
 - overview 9
 - PuTTY 9
- SSPC 7
- status
 - cluster 10
- storage area network (SAN)
 - configuring 46
 - fabric overview 33
- subsystem
 - configuration rules 46
- subsystem device driver (SDD) 5
- switches
 - Brocade 36
 - Cisco 36
 - configuring 35
 - director class 38
 - fibre-channel 35
 - McData 36
 - mixing 36
 - operating over long distances 43
 - redundant ac power 6
 - zoning 41

T

- trademarks 59
- troubleshooting
 - event notification e-mail 9
 - using Assist On-site 8

U

- uninterruptible power supply
 - 2145 uninterruptible power supply environment 19
 - 2145-1U uninterruptible power supply environment 19, 22
 - overview 6
 - power cables 22
 - ports and connections 25
 - preparing environment 21

V

- virtual disks (Vdisks)
 - mirroring 14
- virtual disks (VDisks)
 - modes and migration 16
 - nodes 51
 - space-efficient 14
- virtualization
 - SAN Volume Controller 2

W

- Web browsers
 - requirements 25
 - supported for SAN Volume Controller
 - console 25
- Web sites xv

Z

- zoning
 - controllers 39
 - Global Mirror 43
 - guidelines 39
 - hosts 39
 - Metro Mirror 43
 - overview 41

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

IBM System Storage SAN Volume Controller
Planning Guide
Version 4.3.0

Publication No. GA32-0551-03

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

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

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

Comments:

Thank you for your support.

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

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

Name

Address

Company or Organization

Phone No.

E-mail address



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

GA32-0551-03



Spine information:



IBM System Storage SAN Volume
Controller

SAN Volume Controller Planning Guide

Version 4.3.0